arXiv:1410.0116v1 [math.NA] 1 Oct 2014

A stable, polynomial-time algorithm for the eigenpair problem

Peter Bürgisser*
Technische Universität Berlin
Institut für Mathematik
10623 Berlin, GERMANY
e-mail: pbuerg@math.tu-berlin.de

Felipe Cucker†
Department of Mathematics
City University of Hong Kong
HONG KONG
e-mail: macucker@cityu.edu.hk

Abstract. We describe algorithms for computing eigenpairs (eigenvalue–eigenvector) of a complex $n \times n$ matrix A. These algorithms are numerically stable, strongly accurate, and theoretically efficient (i.e., polynomial-time). We do not believe they outperform in practice the algorithms currently used for this computational problem. The merit of our paper is to give a positive answer to a long-standing open problem in numerical linear algebra.

So the problem of devising an algorithm [for the eigenvalue problem] that is numerically stable and globally (and quickly!) convergent remains open.
J. Demmel [9, page 139]

1 Introduction

1.1 The problem

The quotation from Demmel opening this article, though possibly puzzling for those who day-to-day satisfactorily solve eigenvalue problems, summarizes a long-standing open problem in numerical linear algebra. The first algorithm that comes to mind for computing eigenvalues —to compute the characteristic polynomial χ_A of A and

*Partially funded by DFG research grant BU 1371/2-2
†Partially funded by a GRF grant from the Research Grants Council of the Hong Kong SAR (project number CityU 100810).

then compute (i.e., approximate) its zeros— has proved to be numerically unstable in practice. The so called Wilkinson's polynomial,

$$w(x) := \prod_{i=1}^{20}(x-i) = x^{20} + w_{19}x^{19} + \cdots + w_1 x + w_0$$

is often used to illustrate this fact. For a diagonal matrix D with diagonal entries $1, 2, \ldots, 20$ (and therefore with $\chi_D(x) = w(x)$) an error of 2^{-23} in the computation of $w_{19} = 210$ produces, even if the rest of the computation is done error-free, catastrophic variations in the zeros of χ_D. For instance, the zeros at 18 and 19 collide into a double zero close to 18.62, which will unfold into two complex conjugate zeros if the error in w_{19} is further increased. And yet, there is nothing wrong in the nature of D (in numerical analysis terms, and we will be more detailed below, D is a well-conditioned matrix for the eigenvalue problem). The trouble appears to lie in the method.

Barred from using this immediate algorithm due to its numerical unstability, researchers devoted efforts to come up with alternate methods which would appear to be stable. Among those proposed, the one that is today's algorithm of choice is the iterated QR with shifts. This procedure behaves quite efficiently in general and yet, as Demmel pointed out in 1997 [9, p. 139],

> after more than 30 years of dependable service, convergence failures of this algorithm have quite recently been observed, analyzed, and patched [...]. But there is still no global convergence proof, even though the current algorithm is considered quite reliable.

Our initial quotation followed these words in Demmel's text. It demanded for an algorithm which will be numerically stable and for which, convergence, and if possible small complexity bounds, can be established. Today, 17 years after Demmel's text, this demand retains all of its urge.

The only goal of this paper is to give a positive answer to it.

1.2 A few words on approximations

It must be said upfront that we do not think the algorithm we propose will outperform, in general, iterated QR with shifts. It nonetheless possesses some worthy features which we want to describe in this introduction. The key one, we already mentioned, is that both convergence and complexity bounds can be established for it. It is also numerically stable. But in addition, it is *strongly accurate*.

A starting point to understand the meaning of this last claim, is the observation that there are two different obstructions to the exact computation of an eigenvalue. Firstly, the use of finite precision, and the ensuing errors accumulating during the computational process. The expression *numerically stable* is usually vested on algorithms for which this accumulated error on the computed quantities is not much

larger than that resulting from an error-free computation on an input datum which has been read (and approximated) with machine precision. Secondly, the nonlinear character of the equations defining eigenvalues in terms of the given matrix. For $n \geq 5$, we learned from Galois, we cannot write down these eigenvalues in terms of the matrix' entries, not even using radicals. Hence, we can only compute approximations of them and this is so *even assuming infinite precision in the computation.*

The expression *strongly accurate* refers to the quality of these approximations. It is common to find in the literature (at least) three notions of approximation which we next briefly describe. To simplify, we illustrate with the computation of a value $\zeta \in \mathbb{C}$ from a datum $A \in \mathbb{C}^N$ (and the reader may well suppose that this computation is done with infinite precision). We let $\widetilde{\zeta}$ be the quantity actually computed and we consider the following three requirements on it:

Backward approximation. The element $\widetilde{\zeta}$ is the solution of a datum $\widetilde{A}$ close to A. Given $\varepsilon > 0$, we say that $\widetilde{\zeta}$ is an ε-backward approximation when $\|A - \widetilde{A}\| \leq \varepsilon$.

Forward approximation. The quantity $\widetilde{\zeta}$ is close to ζ. Given $\varepsilon > 0$, we say that $\widetilde{\zeta}$ is an ε-forward approximation when $|\zeta - \widetilde{\zeta}| \leq \varepsilon$.

Approximation à la Smale. An appropriate version of Newton's iteration, starting at $\widetilde{\zeta}$, converges immediately, quadratically fast, to ζ.

These requirements are increasingly demanding. For instance, if ζ is an ε-backward approximation then the forward error $|\zeta - \widetilde{\zeta}|$ is bounded, roughly, by $\varepsilon\,\mathsf{cond}(A)$. Here $\mathsf{cond}(A)$ is the condition number of A, a quantity usually greater than 1. So, in general, ε-backward approximations are not ε-forward approximations, and if A is poorly conditioned $\widetilde{\zeta}$ may be a poor forward approximation of ζ. We also observe that if $\widetilde{\zeta}$ is an approximation à la Smale we can obtain an ε-forward approximation from $\widetilde{\zeta}$ by performing $\mathcal{O}(\log|\log\varepsilon|)$ Newton's steps. Obtaining an approximation à la Smale from an ε-forward approximation is a much less obvious process.

When we say that our algorithm is strongly accurate, we refer to the fact that the returned eigenpairs are approximations à la Smale of true eigenpairs.

1.3 A few words on complexity

The cost, understood as the number of arithmetic operations performed, of computing an approximation of an eigenpair for a matrix $A \in \mathbb{C}^{n\times n}$, depends on the matrix A itself. Actually, and this is a common feature in numerical analysis, it depends on the condition $\mathsf{cond}(A)$ of the matrix A. But this condition is not known a priori. It was therefore advocated by Smale [18] to eliminate this dependency in complexity bounds by endowing data space with a probability distribution and estimating expected costs. This idea has its roots in early work of Goldstine and von Neumann [23].

In our case, data space is $\mathbb{C}^{n\times n}$, and a common probability measure to endow it with is the standard Gaussian. Expectations of cost w.r.t. this measure yield expressions in n usually referred to as *average cost*. A number of considerations, including the suspicion that the use of the standard Gaussian could result in complexity bounds which are too optimistic compared with "real life", prompted Spielman and Teng to introduce a different form of probabilistic analysis, called *smoothed analysis*. In this, one replaces the average analysis' goal of showing that

> for a random A it is unlikely that the cost for A will be large

by the following

> for all $\overline{A}$, it is unlikely that a slight random perturbation $A = \overline{A} + \Delta A$ will require a large cost.

The expectations obtained for a smoothed analysis will now be functions of both the dimension n and some measure of dispersion for the random perturbations (e.g., a variance).

Smoothed analysis was first used for the simplex method of linear programming [21]. Two survey expositions of its rationale are in [20, 22]. One may argue that it has been well accepted by the scientific community from the fact that Spielman and Teng were awarded the Gödel 2008 and Fulkerson 2009 prizes for it (the former by the theoretical computer science community and the latter by the optimization community). Also, in 2010, Spielman was awarded the Nevanlinna prize, and smoothed analysis appears in the laudatio of his work.

In this paper we will exhibit bounds for the cost of our algorithm both for average and smoothed analyses.

1.4 A few words on numerical stability

The algorithm we deal with in this paper belongs to the class of homotopy continuation methods. Experience has shown that algorithms in this class are very stable and stability analyses have been done for some of them, e.g. [5]. Because of this, we will assume infinite precision all along this paper and steer clear of any form of stability analysis. We nonetheless observe that such an analysis can be easily carried out following the steps in [5].

1.5 Previous and related work

Homotopy continuation methods go back, at least, to the work of Lahaye [10]. A detailed survey of their use to solve polynomial equations is in [11].

In the early 1990s Shub and Smale set up a program to understand the cost of solving square systems of complex polynomial equations using homotopy methods. In a collection of articles [12, 13, 14, 15, 16], known as the *Bézout series*, they put

in place many of the notions and techniques that occur in this article. The Bézout series did not, however, conclusively settle the understanding of the cost above, and in 1998 Smale proposed it as the 17th in his list of problems for the mathematicians of the 21st century [19]. The problem is not yet fully solved but significant advances appear in [3, 4, 6].

The results in these papers cannot be directly used for the eigenpair problem since instances of the latter are ill-posed as polynomial systems. But the intervening ideas can be reshaped to attempt a tailor-made analysis for the eigenpair problem. A major step in this direction was done by Armentano in his PhD thesis (see [1]), where a condition number μ for the eigenpair problem was defined and exhaustively studied. A further step was taken in [2] where μ was used to analyze a randomized algorithm for the Hermitian eigenpair problem. A difference between our paper and both [1] and [2] is that in the latter the technical development binds inputs and outputs (eigenvalues) together. We have found more natural to uncouple them.

Our paper follows this stream of research.

1.6 Structure of the exposition

The remaining of this paper is divided into two parts. In the first one, Section 2 below, we introduce all the technical preliminaries, we describe with details the algorithms, and we state our main results (Theorems 2.22 and 2.23). The condition number μ, Newton's method, the notion of approximate eigenpair, and Gaussian distributions are among these technical preliminaries. The second part, which occupies us in Sections 3 to 7, is devoted to proofs.

2 Preliminaries, Basic Ideas, and Main Result

2.1 Spaces and Metric Structures

Let $\mathbb{C}^{n\times n}$ be the set of $n \times n$ complex matrices. We endow this complex linear space with the restriction of the real part of the *Frobenius Hermitian product* $\langle\ ,\ \rangle_F$ given by

$$\langle A, B\rangle_F := \text{trace}\ (B^*A) = \sum_{i,j=1}^{n} a_{ij}\,\overline{b_{ij}},$$

where $A = (a_{ij})$ and $B = (a_{ij})$. The *Frobenius norm* $\|\ \|_F$ on $\mathbb{C}^{n\times n}$ is the norm induced by $\langle\ ,\ \rangle_F$.

On the product vector space $\mathbb{C}^{n\times n} \times \mathbb{C}$ we introduce the canonical Hermitian inner product structure and its associated norm structure and (Euclidean) distance.

The space $\mathbb{C}^n$ is equipped with the canonical Hermitian inner product $\langle\ ,\ \rangle$. We denote by $\mathbb{P}(\mathbb{C}^n)$ its associated projective space. This is a smooth manifold which carries a natural Riemannian metric, namely, the real part of the *Fubini-Study metric*

on $\mathbb{P}(\mathbb{C}^n)$. The Fubini-Study metric is the Hermitian structure on $\mathbb{P}(\mathbb{C}^n)$ given in the following way: for $x \in \mathbb{C}^n$,

$$\langle w, w' \rangle_x := \frac{\langle w, w' \rangle}{\|x\|^2}, \tag{1}$$

for all w, w' in the Hermitian complement $x^\perp$ of x in $\mathbb{C}^n$. We denote by $d_{\mathbb{P}}$ the associated Riemannian distance.

The space $\mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$ is endowed with the Riemannian product structure.

The Hermitian structure in the spaces $\mathbb{C}^{n\times n}$ and $\mathbb{C}^{n\times n} \times \mathbb{C}$ naturally endows them with a notion of *angle*. The Riemannian distances $d_{\mathbb{S}}$ and $d_{\mathbb{P}}$ in the unit sphere $\mathbb{S}(\mathbb{C}^{n\times n})$ and the projective space $\mathbb{P}(\mathbb{C}^n)$, respectively, are given precisely by the angle between its arguments.

In addition, we will consider the following function on $((\mathbb{C}^{n\times n} \setminus \{0\}) \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n))^2$ given by

$$\mathsf{dist}((A, \lambda, v), (A', \lambda', v'))^2 := \left\| \frac{A}{\|A\|_F} - \frac{A'}{\|A'\|_F} \right\|_F^2 + \left| \frac{\lambda}{\|A\|_F} - \frac{\lambda'}{\|A'\|_F} \right|^2 + d_{\mathbb{P}}(v, v')^2.$$

When restricted to matrices with Frobenius norm 1 this defines a distance on $\mathbb{S}(\mathbb{C}^{n\times n}) \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$. In general it is not, but it remains a convenient measure due to the scaling invariance of the eigenpair problem. The space $\mathbb{S}(\mathbb{C}^{n\times n}) \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$ is naturally endowed with a Riemannian distance d_{R} given by replacing $\|A - A'\|_F$ in the definition of dist by the angle $d_{\mathbb{S}}(A, A')$. Since chords are smaller than their subtending angles we trivially have

$$\mathsf{dist}((A, \lambda, v), (A', \lambda', v')) \leq d_{\mathsf{R}}((A, \lambda, v), (A', \lambda', v')) \tag{2}$$

for all $(A, \lambda, v), (A', \lambda', v') \in \mathbb{S}(\mathbb{C}^{n\times n}) \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$.

2.2 The Varieties $\mathcal{V}$, $\mathcal{W}$, Σ' and Σ

We define the *solution variety* for the eigenpair problem as

$$\mathcal{V} := \left\{ (A, \lambda, v) \in \mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n) : \ (A - \lambda \mathsf{Id})v = 0 \right\}.$$

Proposition 2.1 *The solution variety $\mathcal{V}$ is a smooth submanifold of $\mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$, of the same dimension as $\mathbb{C}^{n\times n}$.*

Proof. See [1, Proposition 2.2]. □

The set $\mathcal{V}$ inherits the Riemannian structure of the ambient space. Associated to it there are natural projections:

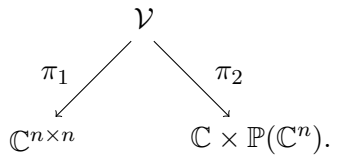

(3)

Because of Proposition 2.1, the derivative $D\pi_1$ at (A,λ,v) is a linear operator between spaces of equal dimension. The *subvariety $\mathcal{W}$ of well-posed triples* is the subset of triples $(A,\lambda,v)\in\mathcal{V}$ for which $D\pi_1(A,\lambda,v)$ is an isomorphism. In particular, when $(A,\lambda,v)\in\mathcal{W}$, the projection π_1 has a branch of its inverse (locally defined) taking $A\in\mathbb{C}^{n\times n}$ to $(A,\lambda,v)\in\mathcal{V}$. This branch of π_1^{-1} is called the *solution map* at (λ,v).

For $v\in\mathbb{P}(\mathbb{C}^n)$ we denote by T_v the tangent space of $\mathbb{P}(\mathbb{C}^n)$ at v. We then have $T_v:=\{x\in\mathbb{C}^n\mid\langle x,w\rangle=0\}$ for any $w\in\mathbb{C}^n$ such that $v=[w]$. Let $P_{v^\perp}:\mathbb{C}^n\to T_v$ be the orthogonal projection. Given $(A,\lambda,v)\in\mathbb{C}^{n\times n}\times\mathbb{C}\times\mathbb{P}(\mathbb{C}^n)$, we let $A_{\lambda,v}:T_v\to T_v$ be the linear operator given by

$$A_{\lambda,v}:=P_{v^\perp}\circ(A-\lambda\mathsf{Id})|_{T_v}. \tag{4}$$

We will prove (in Proposition 6.3 below) that the set of well-posed triples is given by

$$\mathcal{W}=\{(A,\lambda,v)\in\mathcal{V}:\ A_{\lambda,v}\ \text{is invertible}\} \tag{5}$$

(see also Lemma 2.7 in [1]).

We define $\Sigma':=\mathcal{V}\setminus\mathcal{W}$ to be the variety of *ill-posed triples*, and $\Sigma=\pi_1(\Sigma')\subset\mathbb{C}^{n\times n}$ the *discriminant variety*, i.e., the subset of *ill-posed inputs*.

Remark 2.2 From (5) it is clear that the subset Σ' is the set of triples $(A,\lambda,v)\in\mathcal{V}$ such that λ is an eigenvalue of A of algebraic multiplicity at least 2. It follows that Σ is the set of matrices $A\in\mathbb{C}^{n\times n}$ with multiple eigenvalues. In particular, when $A\in\mathbb{C}^{n\times n}\setminus\Sigma$, the eigenvalues of A are pairwise different and $\pi_1^{-1}(A)$ is the set of triples $(A,\lambda_1,v_1),\dots,(A,\lambda_n,v_n)$, where (λ_i,v_i), $i=1,\dots,n$, are the eigenpairs of A.

Proposition 2.3 *The discriminant variety $\Sigma\subset\mathbb{C}^{n\times n}$ is a complex algebraic hypersurface. Consequently, for all $n\geq 2$, we have $\dim_{\mathbb{R}}\Sigma=n^2-2$.*

Proof. See [7, Proposition 20.18]. □

2.3 Unitary invariance

Let $\mathcal{U}(n)$ be the group of $n\times n$ unitary matrices. The group $\mathcal{U}(n)$ naturally acts on $\mathbb{P}(\mathbb{C}^n)$ by $U([w]):=[Uw]$. In addition, $\mathcal{U}(n)$ acts on $\mathbb{C}^{n\times n}$ by conjugation (i.e.,

$U \cdot A := UAU^{-1}$), and on $\mathbb{C}^{n\times n} \times \mathbb{C}$ by $U \cdot (A, \lambda) := (UAU^{-1}, \lambda)$. These actions define an action on the product space $\mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$, namely,

$$U \cdot (A, \lambda, v) := (UAU^{-1}, \lambda, Uv). \tag{6}$$

Remark 2.4 The varieties $\mathcal{V}$, $\mathcal{W}$, Σ', and Σ, are invariant under the action of $\mathcal{U}(n)$ (see [1] for details).

2.4 Condition of a triple

In a nutshell, condition numbers measure the worst possible output error resulting from a small perturbation on the input data. More formally, a condition number is the operator norm of a solution map such as the branches of π^{-1} mentioned in §2.1 above, (see [7, §14.1.2] for a general exposition).

In the case of the eigenpair problem, one can define two condition numbers for eigenvalue and eigenvector, respectively. Armentano has shown, however, that one can merge the two in a single one (see Section 3 in [1] for details). Deviating slightly from [1], we define the *condition number* of $(A, \lambda, v) \in \mathcal{W}$ as

$$\mu(A, \lambda, v) := \|A\|_F \|A^{-1}_{\lambda,v}\|, \tag{7}$$

where $\|\ \|$ is the operator norm.

Remark 2.5 The condition number μ is invariant under the action of the unitary group $\mathcal{U}(n)$, i.e., $\mu(UAU^{-1}, \lambda, Uv) = \mu(A, \lambda, v)$ for all $U \in \mathcal{U}(n)$. Also, it is easy to see, μ is scale invariant on the first two components. That is, $\mu(sA, s\lambda, v) = \mu(A, \lambda, v)$ for all nonzero real s.

Lemma 2.6 (Lemma 3.8 in [1]) *For $(A, \lambda, v) \in \mathcal{V}$ we have $\mu(A, \lambda, v) \geq \frac{1}{\sqrt{2}}$.* □

The essence of condition numbers is that the measure how much may outputs vary when inputs are slightly perturbed. The following result, which we will prove in Section 3, quantifies this property for μ.

Proposition 2.7 *Let $\Gamma : [0, 1] \to \mathcal{V}$, $\Gamma(t) = (A_t, \lambda_t, v_t)$ be a smooth curve such that A_t lies in the unit sphere of $\mathbb{C}^{n\times n}$, for all t. Write $\mu_t := \mu(\Gamma(t))$. Then we have, for all $t \in [0, 1]$,*

$$|\dot{\lambda_t}| \leq \sqrt{1 + \mu_t^2}\ \|\dot{A}_t\|, \qquad \|\dot{v_t}\| \leq \mu_t\ \|\dot{A}_t\|.$$

In particular,

$$\left\|\dot{\Gamma}(t)\right\| \leq \sqrt{6}\ \mu_t\ \|\dot{A}_t\|.$$

Remark 2.8 Since the property of $A_{\lambda,v}$ being invertible is Zariski open on $\mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$, the condition number μ can be naturally extended to a Zariski open neighborhood of $\mathcal{W}$ in $\mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$. We will denote this extension also by μ.

In addition, when $A_{\lambda,v}$ is non-invertible we will let $\mu(A,\lambda,v) := \infty$, so that now μ is well-defined in all of $\mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$.

Condition numbers themselves vary in a controlled manner. The following Lipschitz property makes this statement precise.

Theorem 2.9 *Let $A, A' \in \mathbb{C}^{n\times n}$ be such that $\|A\|_F = \|A'\|_F = 1$, let $v, v' \in \mathbb{C}^n$ be nonzero, and let $\lambda, \lambda' \in \mathbb{C}$ be such that $Av = \lambda v$. Suppose that*

$$\mu(A,\lambda,v)\ \mathsf{dist}((A,\lambda,v),(A',\lambda',v'))\ \leq\ \frac{\varepsilon}{12.5}$$

for $0 < \varepsilon < 0.37$. Then we have

$$\frac{1}{1+\varepsilon}\,\mu(A,\lambda,v)\ \leq\ \mu(A',\lambda',v')\ \leq\ (1+\varepsilon)\,\mu(A,\lambda,v).$$

We give the proof of Theorem 2.9 in Section 4.

Condition numbers are generally associated to input data. In the case of a problem with many possible solutions (of which returning an eigenpair of a given matrix is a clear case) one can derive the condition of a data from a notion of condition for each of these solutions. A discussion of this issue is given in [7, Section 6.8]. For the purposes of this paper, we will be interested in two such derivations: the *maximum condition number* of A,

$$\mu_{\max}(A) := \max_{j\leq n} \mu(A,\lambda_j,v_j),$$

and the *mean square condition number* of A,

$$\mu_{\mathsf{av}}(A) := \left(\frac{1}{n}\sum_{j=1}^{n} \mu^2(A,\lambda_j,v_j)\right)^{\frac{1}{2}}.$$

We close this paragraph observing that restricted to the class of normal matrices, the condition number μ admits the following elegant characterization.

Lemma 2.10 (Lemma 3.12 in [1]) *Let $A \in \mathbb{C}^{n\times n} \setminus \Sigma$ be normal, and let $(\lambda_1,v_1),\ldots,(\lambda_n,v_n)$ be its eigenpairs. Then*

$$\mu(A,\lambda_1,v_1) = \frac{\|A\|_F}{\min_{j=2,\ldots,n} |\lambda_j - \lambda_1|} \qquad \square$$

2.5 Newton's method and approximate eigenpairs

For a nonzero matrix $A \in \mathbb{C}^{n\times n}$, we define the *Newton map* associated to A,

$$N_A : \mathbb{C} \times (\mathbb{C}^n \setminus \{0\}) \to \mathbb{C} \times (\mathbb{C}^n \setminus \{0\}),$$

by $N_A(\lambda, v) = (\lambda - \dot{\lambda}, v - \dot{v})$ where

$$\dot{v} = {A_{\lambda,v}}^{-1} P_{v^\perp}(A - \lambda \,\mathsf{Id})v, \qquad \dot{\lambda} = \frac{\langle (A - \lambda \,\mathsf{Id})(v - \dot{v}), v\rangle}{\langle v, v\rangle}.$$

This map is defined for every $(\lambda, v) \in \mathbb{C} \times (\mathbb{C}^n \setminus \{0\})$ such that $A_{\lambda,v} := P_{v^\perp}(A - \lambda \,\mathsf{Id})|_{T_v}$ is invertible. It was introduced in [1] as the Newton operator associated to the evaluation map $(\lambda, v) \mapsto (A - \lambda \,\mathsf{Id})v$ for a fixed A. See Section 4 of [1] for more details.

Definition 2.11 Given $(A, \lambda, v) \in \mathcal{W}$ we say that $(\zeta, w) \in \mathbb{C} \times (\mathbb{C}^n \setminus \{0\})$ is an *approximate eigenpair* of A with associated eigenpair (λ, v) when for all $k \geq 1$ the kth iterate $N_A^k(\zeta, w)$ of the Newton map at (ζ, w) is well defined and satisfies

$$\mathsf{dist}\big((A, N_A^k(\zeta, w)), (A, \lambda, v)\big) \leq \left(\frac{1}{2}\right)^{2^k-1} \mathsf{dist}\big((A, \zeta, w), (A, \lambda, v)\big).$$

Remark 2.12 Note that, if $N_A(\zeta, w) = (\zeta', w')$ then $N_{sA}(s\zeta, \beta w) = (s\zeta', \beta w')$, for every $s \in \mathbb{C} \setminus \{0\}$ and $\beta \in \mathbb{C} \setminus \{0\}$. Hence, N_A is correctly defined on $\mathbb{C} \times \mathbb{P}(\mathbb{C}^n)$ and the notion of approximate eigenpair scales correctly in the sense that if (λ, v) is an approximate eigenpair of A then $(s\lambda, v)$ is an approximate eigenpair of sA, for all $s \in \mathbb{C} \setminus \{0\}$.

Remark 2.13 The notion of approximate solution as a point where Newton's method converges to a true solution immediately and quadratically fast was introduced by Steve Smale [17]. It allows to elegantly talk about polynomial time without dependencies on pre-established accuracies. In addition, these approximate solutions are "good approximations" (as mentioned in the statement of the Main Theorem) in a very strong sense. The distance to the exact solution dramatically decreases with a single iteration of Newton's method.

The following result estimates, in terms of the condition of an eigenpair, the radius of a ball of approximate eigenpairs associated to it. For a proof see [1, Theorem 5].

Theorem 2.14 *Let $A \in \mathbb{C}^{n\times n}$ with $\|A\|_F = 1$ and $(\lambda, v), (\lambda_0, v_0) \in \mathbb{C} \times (\mathbb{C}^n \setminus \{0\})$. If (λ, v) is a well-posed eigenpair of A and*

$$\mathsf{dist}\big((\lambda, v), (\lambda_0, v_0)\big) < \frac{c_0}{\mu(A, \lambda, v)}$$

then (λ_0, v_0) is an approximate eigenpair of A with associated eigenpair (λ, v). One may choose $c_0 = 0.2881$. □

Remark 2.15 We note that $N_A(\zeta, w)$ can be computed from the matrix A and the pair (ζ, w) in $\mathcal{O}(n^3)$ operations, since the cost of this computation is dominated by that of inverting a matrix.

2.6 Gaussian Measures on $\mathbb{C}^{n\times n}$

Let $\sigma > 0$. We say that the complex random variable $Z = X + \sqrt{-1}Y$ has distribution $\mathcal{N}_{\mathbb{C}}(0, \sigma^2)$ when the real part X and the imaginary part Y are independent and identically distributed (i.i.d.) drawn from $\mathcal{N}(0, \frac{\sigma^2}{2})$, i.e., they are Gaussian centered random variables with variance $\frac{\sigma^2}{2}$.

If $Z \sim \mathcal{N}_{\mathbb{C}}(0, \sigma^2)$ then its density $\varphi\colon \mathbb{C} \to \mathbb{R}$ with respect to the Lebesgue measure is given by

$$\varphi(z) := \frac{1}{\pi\sigma^2} e^{-\frac{|z|^2}{\sigma^2}}.$$

We will write $v \sim \mathcal{N}_{\mathbb{C}}(0, \mathbf{1}_n\sigma^2)$ to indicate that the vector $v \in \mathbb{C}^n$ is random with i.i.d. coordinates drawn from $\mathcal{N}_{\mathbb{C}}(0, \sigma^2)$. Also, we say that $A \in \mathbb{C}^{n\times n}$ is (isotropic) *Gaussian*, if its entries are i.i.d. Gaussian random variables. In this case we write $A \sim \mathcal{N}_{\mathbb{C}}(0, \sigma^2\mathsf{Id}_{n\times n})$, or simply $A \sim \mathcal{N}(0, \sigma^2\mathsf{Id})$ (since we will only deal with square complex matrices here and in general the dimension n will be clear from the context).

If $\overline{A} \in \mathbb{C}^{n\times n}$ and $G \sim \mathcal{N}(0, \sigma^2\mathsf{Id})$, we say that the random matrix $A = G + \overline{A}$ has the *Gaussian distribution centered at* $\overline{A}$, and we write $A \sim \mathcal{N}(\overline{A}, \sigma^2\mathsf{Id})$. The density of this distribution is given by

$$\varphi^{\overline{A},\sigma}_{n\times n}(A) := \frac{1}{(\pi\sigma^2)^{n^2}}\, e^{-\frac{\|A-\overline{A}\|_F^2}{\sigma^2}}.$$

Crucial in our development is the following result giving a bound on average condition for Gaussian matrices arbitrarily centered. We will prove it in Section 6.

Theorem 2.16 *For $\overline{Q} \in \mathbb{C}^{n\times n}$ and $\sigma > 0$ we have*

$$\mathop{\mathbb{E}}_{Q\sim\mathcal{N}(\overline{Q},\sigma^2\mathsf{Id})} \Big(\frac{\mu^2_{\mathsf{av}}(Q)}{\|Q\|^2}\Big) \;\leq\; \frac{en}{2\sigma^2}.$$

Remark 2.17 (i) We note that no bound on the norm of $\overline{Q}$ is required here. Indeed, using $\mu_{\mathsf{av}}(sQ) = \mu_{\mathsf{av}}(Q)$, it is easy to see that the assertion for a pair $(\overline{Q}, \sigma)$ implies the assertion for $(s\overline{Q}, s\sigma)$, for any $s > 0$.

(ii) Because of Proposition 2.3, with probability one, matrices drawn from $\mathcal{N}(\overline{Q}, \sigma^2\mathsf{Id})$ have all its eigenvalues different. Therefore the expected value in Theorem 2.16 is well-defined.

2.7 Truncated Gaussians and smoothed analysis

For $T, \sigma > 0$, we define the *truncated Gaussian* $\mathcal{N}_T(0, \sigma^2 \mathsf{Id})$ on $\mathbb{C}^{n\times n}$ to be the distribution given by the density

$$\rho_T^\sigma(A) = \begin{cases} \frac{\varphi_{n\times n}^{0,\sigma}(A)}{P_{T,\sigma}} & \text{if } \|A\|_F \leq T, \\ 0 & \text{otherwise,} \end{cases} \tag{8}$$

where $P_{T,\sigma} := \mathsf{Prob}_{f\sim\mathcal{N}(0,\sigma^2\mathsf{Id})}\{\|f\| \leq T\}$, and, we recall, $\varphi_{n\times n}^{0,\sigma}$ is the density of $\mathcal{N}(0, \sigma^2\mathsf{Id})$. For the rest of this paper we fix the threshold $T := \sqrt{2}\, n$. The fact that $\|A\|_F^2$ is chi-square distributed with $2n^2$ degrees of freedom, along with [8, Corollary 6] yield the following result.

Lemma 2.18 *We have $P_{T,\sigma} \geq \frac{1}{2}$ for all $0 < \sigma \leq 1$.* □

The space $\mathbb{C}^{n\times n}$ of matrices with the Frobenius norm and the space $\mathbb{C}^{n^2}$ with the canonical Hermitian product are isomorphic as Hermitian product spaces. Hence, the Gaussian $\mathcal{N}(0, \sigma^2\mathsf{Id}_{n\times n})$ on the former corresponds to the Gaussian $\mathcal{N}(0, \sigma^2\mathbf{1}_{n^2})$ on the latter, and we can deduce invariance of $\mathcal{N}(0, \sigma^2\mathsf{Id})$ under the action of $\mathcal{U}(n^2)$ (in addition to that for conjugation under $\mathcal{U}(n)$ discussed in §2.3). The same is true for the truncated Gaussian $\mathcal{N}_T(0, \sigma^2\mathsf{Id})$. In particular, the pushforward of both distributions for the projection $\mathbb{C}^{n\times n} \setminus \{0\} \to \mathbb{S}(\mathbb{C}^{n\times n})$, $A \mapsto \frac{A}{\|A\|_F}$, is the uniform distribution $\mathscr{U}(\mathbb{S}(\mathbb{C}^{n\times n}))$ (see [7, Chapter 2] for details) and we have

$$\mathop{\mathbb{E}}_{A\sim\mathcal{N}(0,\sigma^2\mathsf{Id})} F(A) = \mathop{\mathbb{E}}_{A\sim\mathcal{N}_T(0,\sigma^2\mathsf{Id})} F(A) = \mathop{\mathbb{E}}_{A\sim\mathscr{U}(\mathbb{S}(\mathbb{C}^{n\times n}))} F(A). \tag{9}$$

for any integrable scale invariant function $F : \mathbb{C}^{n\times n} \to \mathbb{R}$.

Complexity analysis has traditionally been carried out either in the *worst-case* or in an *average-case*. More generally, for a function $F : \mathbb{R}^m \to \mathbb{R}_+$, the former amounts to the evaluation of $\sup_{a\in\mathbb{R}^m} F(a)$ and the latter to that of $\mathbb{E}_{a\sim\mathscr{D}} F(a)$ for some probability distribution $\mathscr{D}$ on $\mathbb{R}^m$. Usually, $\mathscr{D}$ is taken to be the standard Gaussian $\mathcal{N}(0, \mathsf{Id})$. With the beginning of the century, Daniel Spielman and Shang-Hua Teng introduced a third form of analysis, *smoothed analysis*, which is meant to interpolate between worst-case and average-case. We won't elaborate here on the virtues of smoothed analysis; a defense of these virtues can be found, e.g., in [20, 22] or in [7, §2.2.7]. We will instead limit ourselves to the description of what smoothed analysis is and which form it will take in this paper.

The idea is to replace the two operators above (supremum and expectation) by a combination of the two, namely,

$$\sup_{\overline{a}\in\mathbb{R}^m} \mathop{\mathbb{E}}_{a\sim\mathscr{D}(\overline{a},\sigma)} F(a)$$

where $\mathscr{D}(\overline{a}, \sigma)$ is a distribution "centered" at $\overline{a}$ having σ as a measure of dispersion. A typical example is the Gaussian $\mathcal{N}(\overline{a}, \sigma^2 \mathsf{Id})$. Another example, used for scale invariant functions F, is the uniform measure on a spherical cap centered at $\overline{a}$ and with angular radius σ on the unit sphere $\mathbb{S}(\mathbb{R}^m)$ (reference [7] exhibits smoothed analyses for both examples of distribution). In this paper we will perform a smoothed analysis with respect to a truncated Gaussian. More precisely, we will be interested in quantities of the form

$$\sup_{\overline{A} \in \mathbb{C}^{n \times n}} \mathop{\mathbb{E}}_{A \sim \mathcal{N}_T(\overline{A}, \sigma^2 \mathsf{Id})} F(A)$$

where F will be a measure of computational cost for the eigenpair problem. We note that, in addition to the usual dependence on n, this quantity depends also on σ and tends to ∞ when σ tends to 0. When F is scale invariant, as in the case of μ_{av} or $\mu_{\max}$, it is customary to restrict attention to matrices of norm 1. That is, to study the following quantity:

$$\sup_{\overline{A} \in \mathbb{S}(\mathbb{C}^{n \times n})} \mathop{\mathbb{E}}_{A \sim \mathcal{N}_T(\overline{A}, \sigma^2 \mathsf{Id})} F(A). \tag{10}$$

2.8 The eigenpair continuation algorithm

We are ready to describe the main algorithmic construct in this paper. For the algorithmic purposes, it will be more convenient to view the solution variety as the corresponding subset of $\mathbb{C}^{n \times n} \times \mathbb{C} \times (\mathbb{C}^n \setminus \{0\})$, which, abusing notation, we still denote by $\mathcal{V}$.

Suppose that we are given an input matrix $A \in \mathbb{C}^{n \times n}$ and an *initial triple* (M, λ, v) in the solution variety $\mathcal{V} \subseteq \mathbb{C}^{n \times n} \times \mathbb{C} \times (\mathbb{C}^n \setminus \{0\})$ such that A and M are $\mathbb{R}$-linearly independent. Let $\alpha := d_{\mathbb{S}}(M, A) \in (0, \pi)$ denote the *angle* between the rays $\mathbb{R}_+ A$ and $\mathbb{R}_+ M$. Consider the line segment $[M, A]$ in $\mathbb{C}^{n \times n}$ with endpoints M and A.

We parameterize this segment by writing

$$[M, A] = \{Q_\tau \in \mathbb{C}^{n \times n} \mid \tau \in [0, 1]\}$$

with Q_τ being the only point in $[M, A]$ such that $d_{\mathbb{S}}(M, Q_\tau) = \tau \alpha$ (see Figure 1).

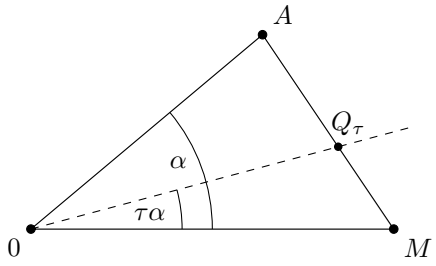

Figure 1: The family Q_τ, $\tau \in [0, 1]$.

If the line segment $[M, A]$ does not intersect the discriminant variety Σ, then starting at the eigenpair (λ, v) of M, the map $[0,1] \to \mathbb{C}^{n\times n}, \tau \mapsto Q_\tau$, can be uniquely extended to a continuous map

$$[0,1] \to \mathcal{V}, \quad \tau \mapsto (Q_\tau, \lambda_\tau, v_\tau), \tag{11}$$

such that $(\lambda_0, v_0) = (\lambda, v)$. We call this map the *lifting* of $[M, A]$ with origin (M, λ, v). We shall also call $\tau \mapsto (Q_\tau, \lambda_\tau, v_\tau)$ the *solution path* in $\mathcal{V}$ corresponding to the input matrix A and initial triple (M, λ, v).

Our algorithm relies on the obvious fact that the pair (λ_1, v_1), corresponding to $\tau = 1$, is an eigenpair of A. We therefore want to find an approximation of this pair and to do so, the idea is to start with the eigenpair $(\lambda, v) = (\lambda_0, v_0)$ of $M = Q_0$ and numerically follow the path $(Q_\tau, \lambda_\tau, v_\tau)$. This is done by subdividing the interval $[0,1]$ into subintervals with extremities at $0 = \tau_0 < \tau_1 < \cdots < \tau_K = 1$ and by successively computing approximations (ζ_i, w_i) of $(\lambda_{\tau_i}, v_{\tau_i})$ by Newton's method. To ensure that these are good approximations, we actually want to ensure that for all i, (ζ_i, w_i) is an approximate eigenpair of $Q_{\tau_{i+1}}$. Figure 2 attempts to convey the general idea.

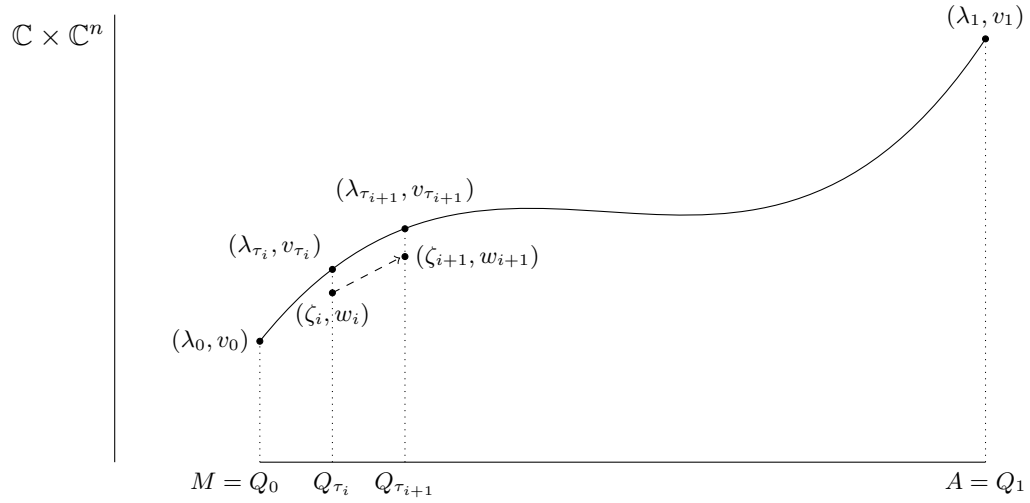

Figure 2: The continuation of the solution path.

The following pseudocode gives a precise description of how this is done. The letter ξ denotes a constant, namely $\xi = 0.001461$.

Algorithm 1 Path-follow

Input: $A, M \in \mathbb{C}^{n\times n}$ and $(\lambda, v) \in \mathbb{C}\times\mathbb{C}^n$

Preconditions: $(M-\lambda\,\mathsf{Id})v = 0$, $M \notin \mathbb{R}A$, $v \neq 0$

$\alpha := d_{\mathbb{S}}(M, A)$, $r := \|A\|_F$, $s := \|M\|_F$

$\tau := 0$, $Q := M$, $(\zeta, w) := (\lambda, v)$

`repeat`

$\Delta\tau := \frac{\xi}{\alpha\mu^2(Q,\zeta,w)}$

$\tau := \min\{1, \tau + \Delta\tau\}$

$t := \frac{s}{s+r(\sin\alpha\cot(\tau\alpha)-\cos\alpha)}$

$Q := tA + (1-t)M$

$(\zeta, w) := N_Q(\zeta, w)$

`until` $\tau = 1$

`return` (ζ, w)

Output: $(\zeta, w) \in \mathbb{C}\times\mathbb{C}^n$

Postconditions: The algorithm halts if the lifting of $[M, A]$ at (λ, v) does not cut Σ'. In this case, (ζ, w) is an approximate eigenpair of A.

Our next result estimates the number of iterations performed by algorithm Path-follow. We summarize its proof in Section 5.

Proposition 2.19 *Suppose that $[M, A]$ does not intersect the discriminant variety Σ. Then the algorithm* Path-follow *stops after at most $K := K(A, M, \lambda)$ steps with*

$$K \;\leq\; 1077\; d_{\mathbb{S}}(M, A) \int_0^1 \mu^2(Q_\tau, \lambda_\tau, v_\tau)\, d\tau.$$

The returned pair (ζ, w) is an approximate eigenpair of A with associated eigenpair (λ_1, v_1). Furthermore, the bound above is optimal up to a constant: we have

$$K \;\geq\; 434\; d_{\mathbb{S}}(M, A) \int_0^1 \mu^2(Q_\tau, \lambda_\tau, v_\tau)\, d\tau.$$

2.9 Initial triples and global algorithms

The Path-follow routine assumes an initial triple (M, λ, v) at hand. We next describe a construction for such initial triples.

For $k \in \mathbb{N}$ we consider the set of points

$$S_k = \left\{ \left(-1 + \frac{2p}{k}, -1 + \frac{2q}{k} \right) \mid 0 \leq p, q \leq k \right\} \subset \mathbb{C}.$$

This is a set of $(k+1)^2$ points which are equidistributed on the square of side length 2, inscribed in the circle $\{z \in \mathbb{C} \mid |z| \leq \sqrt{2}\}$.

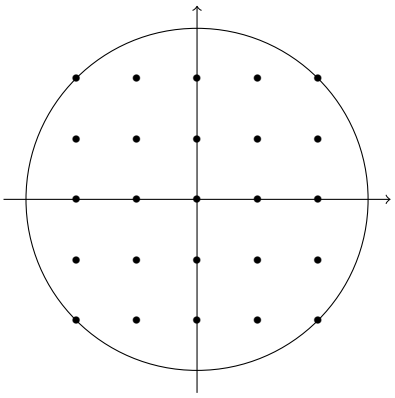

Figure 3: The set S_2.

For $n \geq 2$ we let $k := \lceil \sqrt{n} \rceil - 1$ and define the matrix D_n to be the diagonal matrix whose diagonal entries $z_1, \ldots, z_n$ are the first n elements in S_k (say for the lexicographical ordering). The eigenpairs of D_n are the pairs (z_j, e_j) for $j = 1, \ldots, n$.

Lemma 2.20 *For all $j \leq n$ we have $\mu(D_n, z_j, e_j) = \Theta(n)$.*

PROOF. Let $k := \lceil \sqrt{n} \rceil - 1$. Assume first that k is even and write $z_{pq} := \left(-1 + \frac{2p}{k}, -1 + \frac{2q}{k} \right)$. Then

$$\begin{aligned}
\sum_{p,q=0}^{k} |z_{pq}|^2 &\leq 8 \sum_{p=0}^{\frac{k}{2}} \sum_{q=0}^{p} |z_{pq}|^2 = 8 \sum_{p=0}^{\frac{k}{2}} \sum_{q=0}^{p} \frac{p^2+q^2}{k^2} \\
&= \frac{8}{k^2} \sum_{p=0}^{\frac{k}{2}} \left(p^2(p+1) + \sum_{q=0}^{p} q^2 \right) = \frac{8}{k^2} \sum_{p=0}^{\frac{k}{2}} \left(p^2(p+1) + \frac{p^3}{3} + \frac{p^2}{2} + \frac{p}{6} \right) \\
&= \frac{8}{k^2} \sum_{p=0}^{\frac{k}{2}} \left(\frac{4p^3}{3} + \frac{3p^2}{2} + \frac{p}{6} \right) = \Theta(k^2) = \Theta(n).
\end{aligned}$$

One can similarly prove that the same bound holds for k odd.

It follows that for $n \geq 2$

$$\|D_n\|_F = \left(\sum_{p,q=0}^{\lceil \sqrt{n} \rceil - 1} |z_{pq}|^2 \right)^{1/2} = \Theta(\sqrt{n}).$$

But clearly, the smallest distance between two different eigenvalues of D_n is $\frac{2}{k} = \Theta(n^{-\frac{1}{2}})$. And D_n is diagonal, hence normal. We conclude with Lemma 2.10 that $\mu(D_n) = \Theta(n)$ as claimed. □

We now put together the continuation algorithm and a specific initial triple.

Algorithm 2 Single_Eigenpair

Input: $A \in \mathbb{C}^{n\times n}$

`compute` D_n
`set` $M := \frac{D_n}{\|D_n\|_F}$
`randomly choose` $j \in \{1, \ldots, n\}$
$(\zeta, w) :=$ Path-follow(A, M, m_{jj}, e_j)

Output: $(\zeta, w) \in \mathbb{C} \times \mathbb{C}^n$

Postconditions: The algorithm halts if $[M, A] \cap \Sigma = \emptyset$. In this case, the pair (ζ, w) is an approximate eigenpair of A.

Remark 2.21 The fact that the real codimension of Σ in $\mathbb{C}^{n\times n}$ is 2 (shown in Proposition 2.3) ensures that, almost surely, the segment $[M, A]$ does not intersect Σ and therefore, that almost surely Single_Eigenpair halts.

Given a matrix $A \in \mathbb{C}^{n\times n}$, the cost of algorithm Single_Eigenpair with input A depends on the triple (M, m_{jj}, e_j) which is random in the measure that j is. We therefore consider the *randomized cost* of this algorithm on input A. This amounts to the expected number of iterations of algorithm Path-follow with input A times the $\mathcal{O}(n^3)$ cost of each iteration. The former is given by

$$\mathsf{Num_Iter}(A) := \frac{1}{n}\sum_{j=1}^{n} K(A, M, m_{jj}, e_j).$$

Since we are interested in the *average complexity* of Single_Eigenpair we will further take the expectation of Num_Iter(A) when A is drawn from $\mathcal{N}(0, \sigma^2\mathsf{Id})$. We therefore consider

$$\mathsf{Avg_Num_Iter}(n) := \mathop{\mathbb{E}}_{A\sim\mathcal{N}(0,\sigma^2\mathsf{Id})} \frac{1}{n}\sum_{j=1}^{n} K(A, M, m_{jj}, e_j).$$

Multiplying this expression by the cost $\mathcal{O}(n^3)$ of each iteration, we obtain the *average cost* Avg_Cost(n) of Single_Eigenpair.

We can also consider the smoothed cost of Single_Eigenpair by drawing the input matrix A from $\mathcal{N}_T(\overline{A}, \sigma^2 \mathsf{Id})$ where $\overline{A} \in \mathbb{S}(\mathbb{C}^{n\times n})$ is arbitrary. We thus define

$$\mathsf{Smoothed_Num_Iter}(n) := \sup_{\overline{A}\in\mathbb{S}(\mathbb{C}^{n\times n})} \mathop{\mathbb{E}}_{A\sim\mathcal{N}_T(\overline{A},\sigma^2\mathsf{Id})} \frac{1}{n}\sum_{j=1}^{n} K(A, M, m_{jj}, e_j)$$

and multiplying by $\mathcal{O}(n^3)$ we obtain a corresponding notion of smoothed cost $\mathsf{Smoothed_Cost}(n)$.

We can now state our first main result.

Theorem 2.22 *Algorithm* Single_Eigenpair *returns (almost surely) an approximate eigenpair of its input $A \in \mathbb{C}^{n\times n}$. Its average cost satisfies*

$$\mathsf{Avg_Cost}(n) = \mathcal{O}(n^8).$$

For every $0 < \sigma \leq 1$, its smoothed cost satisfies

$$\mathsf{Smoothed_Cost}(n) = \mathcal{O}\Big(\frac{n^8}{\sigma^2}\Big).$$

We can easily modify algorithm Single_Eigenpair to compute all the eigenpairs.

Algorithm 3 All_Eigenpairs

Input: $A \in \mathbb{C}^{n\times n}$

`compute` D_n
`set` $M := \frac{D_n}{\|D_n\|_F}$
`for` $j \in \{1, \ldots, n\}$ `do`
    $(\zeta_j, w_j) := \mathsf{Path\text{-}follow}(A, M, m_{jj}, e_j)$

Output: $\{(\zeta_1, w_1), \ldots, (\zeta_n, w_n)\} \in (\mathbb{C} \times \mathbb{C}^n)^n$
Postconditions: The algorithm halts if $[M, A] \cap \Sigma = \emptyset$. In this case, the pairs (ζ_j, w_j) are approximate eigenpairs of A with pairwise different associated eigenpairs.

This is no longer a randomized algorithm. In particular, the number of iterations performed by All_Eigenpairs on input A, which is now

$$\mathsf{Num_Iter}(A) := \sum_{j=1}^{n} K(A, M, m_{jj}, e_j),$$

is no longer a random variable. We derive from these quantities the corresponding notions of $\mathsf{Avg_Cost}(n)$ and $\mathsf{Smoothed_Cost}(n)$, for which we state our second main result.

Theorem 2.23 *Algorithm* All_Eigenpairs *returns (almost surely) n approximate eigenpairs of its input $A \in \mathbb{C}^{n\times n}$, with pairwise different associate eigenpairs. Its average cost satisfies*

$$\mathsf{Avg_Cost}(n) = \mathcal{O}(n^9).$$

For every $\sigma \leq 1$ its smoothed cost satisfies

$$\mathsf{Smoothed_Cost}(n) = \mathcal{O}\Big(\frac{n^9}{\sigma^2}\Big).$$

Proof. See Section 7. □

3 Some properties of the condition number μ

There is a general geometric framework for defining condition numbers, see [7, §14.3]. In our situation, this framework takes the following form.

If $(A, \lambda, v) \in \mathcal{V}$ is well-posed, then the projection $\pi_1 \colon \mathcal{V} \to \mathbb{C}^{n\times n}$ (cf. (2)), around (A, λ, v), has a local inverse $U \to \mathcal{V}, A \mapsto (A, G(A))$, that is defined on an open neighborhood U of A in $\mathbb{C}^{n\times n}$. We call G the *solution map*. The map G decomposes as $G = (G_\lambda, G_v)$ where

$$G_\lambda : \mathbb{C}^{n\times n} \to \mathbb{C} \qquad \text{and} \qquad G_v : \mathbb{C}^{n\times n} \to \mathbb{P}(\mathbb{C}^n)$$

associate to matrices B in U an eigenvalue and corresponding eigenvector, respectively. Let

$$DG_\lambda : \mathbb{C}^{n\times n} \to \mathbb{C} \qquad \text{and} \qquad DG_v : \mathbb{C}^{n\times n} \to T_v$$

be the derivatives of these maps at A (which are linear maps between tangent spaces). The condition numbers for the eigenvalue λ and the eigenvector v of A are defined as follows:

$$\mu_\lambda(A, \lambda, v) := \|DG_\lambda\| \qquad \text{and} \qquad \mu_v(A, \lambda, v) := \|DG_v\|$$

where the norms are the operator norms with respect to the chosen norms (on $\mathbb{C}^{n\times n}$ we use the Frobenius norm and on T_v the norm given by (1)). The following result, Lemma 14.17 in [7], gives explicit descriptions of DG_λ and DG_v. Before stating it, we recall that if λ is an eigenvalue of A there exists $u \in \mathbb{C}^n$ such $(A^* - \bar{\lambda}\mathsf{Id})u = 0$. We say that u is a *left eigenvector* of A. Recall the linear map $A_{\lambda,v} \colon T_v \to T_v$ introduced in (4).

Lemma 3.1 *Let $(A, \lambda, v) \in \mathcal{V}$ and let u be a left eigenvector of A with eigenvalue λ. Then:*

(a) *We have $\langle v, u\rangle \neq 0$.*

(b) *If λ is a simple eigenvalue of $A \in \mathbb{C}^{n\times n}$ with right eigenvector v and left eigenvector u, then the derivative of the solution map is given by $DG(A)(\dot{A}) = (\dot{\lambda}, \dot{v})$, where*

$$\dot{\lambda} = \frac{\langle \dot{A}v, u\rangle}{\langle v, u\rangle}, \quad \dot{v} = A_{\lambda,v}^{-1} P_{v^\perp} \dot{A}v. \qquad \square$$

Lemma 3.1 can be used to bound eigenvalue and eigenvector condition numbers. The following result is essentially Prop. 14.15 in [7] (the only difference being that here we use Frobenius norms).

Proposition 3.2 *Choosing the Frobenius norm on $T_A\mathbb{C}^{n\times n} = \mathbb{C}^{n\times n}$ and $\frac{1}{\|v\|}\,\|\ \|$ on T_v, the condition numbers μ_v for the eigenvector problem and μ_λ for the eigenvalue problem satisfy:*

$$\begin{aligned} \mu_\lambda(A,\lambda,v) &= \|DG_\lambda(A)\| &&= \frac{\|u\|\|v\|}{|\langle u, v\rangle|}, \\ \mu_v(A,\lambda,v) &= \|DG_v(A)\| &&= \big\|A_{\lambda,v}^{-1}\big\|. \end{aligned}$$

Proof. For all $\dot{A} \in \mathbb{C}^{n\times n}$ we have, by the Cauchy-Schwarz inequality,

$$|\langle \dot{A}v, u\rangle| \le \|\dot{A}v\|\|u\| \le \|\dot{A}\|_F\|v\|\|u\|.$$

This implies with Lemma 3.1 that

$$\|DG_\lambda(A)\| = \max_{\|\dot{A}\|_F=1} \frac{|\langle \dot{A}v, u\rangle|}{|\langle v, u\rangle|} \le \frac{\|v\|\|u\|}{|\langle v, u\rangle|}.$$

Moreover, there exists a rank one matrix $\dot{A}$ with $\|\dot{A}\| = \|\dot{A}\|_F = 1$ such that $\dot{A}v/\|v\| = u/\|u\|$, cf. [7, Lemma 1.2]. For this choice of $\dot{A}$ we have equality above. This proves the first assertion.

In order to bound $\|DG_v(A)\|$ we note that for all $\dot{A}$,

$$\big\|P_{v^\perp}\dot{A}v\big\| \ \le\ \big\|P_{v^\perp}\big\|\|\dot{A}\|\|v\| \ \le\ \|\dot{A}\|_F\|v\|.$$

Therefore,

$$\max_{\|\dot{A}\|_F=1} \big\|A_{\lambda,v}^{-1}P_{v^\perp}\dot{A}v\big\| \ \le\ \big\|A_{\lambda,v}^{-1}\big\| \max_{\|\dot{A}\|_F=1} \big\|P_{v^\perp}\dot{A}v\big\| \ \le\ \big\|A_{\lambda,v}^{-1}\big\|\,\|v\|.$$

The inequality in the second assertion follows with Lemma 3.1 (and the choice of norm $\|v\|^{-1}\|\ \|$ on T_v). For the equality, let $w \in T_v$ be such that $\|w\| = \|v\|$. Again, by [7, Lemma 1.2], there exists $\dot{A}$ such that $\|\dot{A}\|_F = 1$ and $\dot{A}v = w$, hence $P_{v^\perp}\dot{A}v = w$. This implies

$$\max_{\|\dot{A}\|_F=1} \|A_{\lambda,v}^{-1}\,P_{v^\perp}\dot{A}v\| \ge \max_{\substack{w\in T_v\\ \|w\|=\|v\|}} \|A_{\lambda,v}^{-1}\,w\| = \|v\|\,\|A_{\lambda,v}^{-1}\|,$$

which completes the proof. □

An immediate consequence of Proposition 3.2 is that $\mu_v(A,\lambda,v) = \frac{\mu(A,\lambda,v)}{\|A\|_F}$. We next show that $\mu_\lambda(A,\lambda,v)$ can be similarly bounded in terms of $\mu(A,\lambda,v)$.

Lemma 3.3 *If λ is a simple eigenvalue of A with left eigenvector v and right eigenvector u, then*

$$\mu_\lambda(A,\lambda,v) \;\leq\; \sqrt{1+\mu(A,\lambda,v)^2}.$$

Proof. By unitary invariance, we may assume without loss of generality that $A = \begin{bmatrix}\lambda & a\\ 0 & B\end{bmatrix}$ and $v = (1,0)$, where $a \in \mathbb{C}^{n-1}$ and $B \in \mathbb{C}^{(n-1)\times(n-1)}$. Clearly, $\|a\| \leq \|A\|_F$.

We define $x := ((\lambda \mathsf{Id} - B)^*)^{-1}a^*$. Then $a^* + B^*x = \bar{\lambda}x$, which implies

$$\begin{bmatrix}\bar{\lambda} & 0\\ a^* & B^*\end{bmatrix}\begin{bmatrix}1\\ x\end{bmatrix} = \bar{\lambda}\begin{bmatrix}1\\ x\end{bmatrix}.$$

Hence $u := (1,x)$ is a corresponding right eigenvector of λ. Also,

$$\|x\| \leq \|((\lambda \mathsf{Id} - B)^*)^{-1}\| \cdot \|a^*\| = \|(\lambda \mathsf{Id} - B)^{-1}\| \cdot \|a\| \leq \mu(A,\lambda,v)$$

since $A_{\lambda,v} = B - \lambda \mathsf{Id}$ and $\|a\| \leq \|A\|_F$. Now,

$$\frac{\|u\|\|v\|}{|\langle u,v\rangle|} = \|u\| = \sqrt{1+\|x\|^2}$$

and the assertion follows with Proposition 3.2. □

Proof of Proposition 2.7. The first two inequalities are immediate from Lemma 3.3 and $\mu_v(A,\lambda,v) = \frac{\mu(A,\lambda,v)}{\|A\|_F}$. For the third we have

$$\|\dot{\Gamma}(t)\| \;=\; \|(\dot{A},\dot{\lambda},\dot{v})\| \;\leq\; \|\dot{A}\|\sqrt{1+\mu_t^2+(1+\mu_t^2)} \;\leq\; \|\dot{A}\|\sqrt{6\mu_t^2}$$

the last inequality since $\mu_t \geq \frac{1}{\sqrt{2}}$ (Lemma 2.6). □

4 The Lipschitz property for the eigenpair problem

In this section we prove Theorem 2.9. Recall, we assume that (A,λ,v) is in the solution variety, i.e., $Av = \lambda v$, but we *do not* require that $A'v' = \lambda'v'$. The following result is the main stepping stone.

Proposition 4.1 *Let $A, A' \in \mathbb{C}^{n\times n}$ be such that $\|A\|_F = \|A'\|_F = 1$, let $v, v' \in \mathbb{C}^n$ be nonzero, and let $\lambda, \lambda' \in \mathbb{C}$ be such that $Av = \lambda v$. Suppose that*

$$\mu(A,\lambda,v)\big(\|A'-A\|_F + |\lambda'-\lambda| + d_{\mathbb{P}}(v',v)\big) \;\leq\; \frac{\varepsilon}{7.2}$$

for $0 < \varepsilon \leq 0.37$. Then we have

$$\frac{1}{1+\varepsilon}\,\mu(A,\lambda,v) \;\leq\; \mu(A',\lambda',v') \;\leq\; (1+\varepsilon)\,\mu(A,\lambda,v).$$

Before proceeding, we note that Theorem 2.9 is an immediate consequence of this proposition, since

$$\mathsf{dist}((A,\lambda,v),(A',\lambda',v')) \leq \frac{1}{\sqrt{3}}\Big(\|A'-A\|_F + |\lambda'-\lambda| + d_{\mathbb{P}}(v',v)\Big).$$

We shall provide the proof of Proposition 4.1 in several steps. We begin with the following easy observation whose proof is left to the reader (cf. [7, Lemma 16.40]).

Lemma 4.2 *Suppose that $v', v \in \mathbb{C}^n$ are nonzero with $\delta := d_{\mathbb{P}}(v',v) < \pi/2$. Then the restriction $P_{v^\perp}|_{T_{v'}} : T_{v'} \to T_v$ of $P_{v^\perp}$ to $T_{v'}$ is invertible and we have*

$$\|P_{v^\perp}|_{T_{v'}}\| = \cos\delta, \quad \|(P_{v^\perp}|_{T_{v'}})^{-1}\| = (\cos\delta)^{-1}. \qquad \square$$

In a first step, we fix v and only perturb A and λ.

Lemma 4.3 *Suppose that $A, A' \in \mathbb{C}^{n\times n}$ are such that $\|A\|_F = \|A'\|_F = 1$, $\lambda, \lambda' \in \mathbb{C}$ and $v \in \mathbb{C}^n$ is nonzero. Then,*

$$\mu(A,\lambda,v)\big(\|A'-A\|_F + |\lambda'-\lambda|\big) \;\leq\; \varepsilon < 1$$

implies that

$$(1-\varepsilon)\,\mu(A,\lambda,v) \;\leq\; \mu(A',\lambda',v) \;\leq\; \frac{1}{1-\varepsilon}\,\mu(A,\lambda,v).$$

Proof. Recall that $A_{\lambda,v} = P_{v^\perp}(A - \lambda\mathsf{Id})$, seen as a linear endomorphism of T_v. We have $A'_{\lambda',v} = A_{\lambda,v} + \Delta$, where $\Delta := P_{v^\perp}((A-A') + (\lambda'-\lambda)\mathsf{Id})$, interpreted as a linear endomorphism of T_v. Note that

$$\|\Delta\| \;\leq\; \|P_{v^\perp}\|\cdot\big(\|A-A'\| + \|(\lambda'-\lambda)\mathsf{Id}\|\big) \;\leq\; \|A-A'\| + |\lambda'-\lambda| \;\leq\; \|A-A'\|_F + |\lambda'-\lambda|.$$

Our assumption implies $\|(A_{\lambda,v})^{-1}\|\cdot\|\Delta\| \leq \varepsilon < 1$ (note that $\mu(A,\lambda,v) = \|(A_{\lambda,v})^{-1}\|$ since we assume that $\|A\|_F = 1$). Also, Lemma 15.7 in [7] implies that

$$\|(A'_{\lambda',v})^{-1}\| \;\leq\; \|(A_{\lambda,v})^{-1}\| \cdot \frac{1}{1 - \|(A_{\lambda,v})^{-1}\|\cdot\|\Delta\|}.$$

Using again $\|A\|_F = \|A'\|_F = 1$ it follows that

$$\mu(A', \lambda', v) \;\leq\; \frac{1}{1-\varepsilon}\,\mu(A, \lambda, v). \tag{12}$$

It now suffices to prove that

$$(1-\varepsilon)\,\mu(A, \lambda, v) \;\leq\; \mu(A', \lambda', v).$$

Since this is trivial in the case $\mu(A, \lambda, v) \leq \mu(A', \lambda', v)$, we may assume that $\mu(A', \lambda', v) < \mu(A, \lambda, v)$. Then, by our assumption,

$$\mu(A', \lambda', v)\big(\|A' - A\|_F + |\lambda' - \lambda|\big) \;\leq\; \mu(A, \lambda, v)\big(\|A' - A\|_F + |\lambda' - \lambda|\big) \leq \varepsilon < 1.$$

Hence we can apply (12) by switching the roles of (A', λ', v) and (A, λ, v). This gives

$$\mu(A, \lambda, v) \;\leq\; \frac{1}{1-\varepsilon}\,\mu(A', \lambda', v),$$

which completes the proof. □

In a second step, we fix A and λ and only perturb v. Now we need to assume one of the triples in $\mathcal{V}$.

Lemma 4.4 *Let $A \in \mathbb{C}^{n\times n}$ with $\|A\|_F = 1$, $\lambda \in \mathbb{C}^n$ and $v, v' \in \mathbb{C}$ be nonzero such that $Av = \lambda v$. Then*

$$3.6\,\mu(A, \lambda, v)\,d_{\mathbb{P}}(v, v') \;\leq\; \varepsilon < 1$$

implies that

$$\frac{\mu(A, \lambda, v)}{1+\varepsilon} \;\leq\; \mu(A, \lambda, v') \;\leq\; \frac{\mu(A, \lambda, v)}{1-\varepsilon}.$$

Proof. Since $\mu(A, \lambda, v) < \infty$, λ is a simple eigenvalue of A. Let u denote a corresponding right eigenvector of A, that is, $A^*u = \bar{\lambda}u$. We have

$$\langle u, (\lambda \mathsf{Id} - A)x\rangle = \langle u, \lambda x\rangle - \langle u, Ax\rangle = \bar{\lambda}\langle u, x\rangle - \langle A^*u, x\rangle = 0,$$

hence the image of $\lambda\mathsf{Id} - A$ is contained in T_u and thus equals T_u for reasons of dimension.

Let $\pi_v \colon T_u \to T_v$ denote the restriction to T_u of the orthogonal projection $P_{v^\perp}$. Since $\langle u, v\rangle \neq 0$, the map π_v is an isomorphism. We denote by $\pi_{v'} \colon T_u \to T_{v'}$ the restriction to T_u of the orthogonal projection $P_{v'^\perp}$ and we define $\gamma := \pi_{v'} \circ \pi_v^{-1}$. Moreover, we let $\pi \colon T_{v'} \to T_v$ denote the restriction to $T_{v'}$ of the orthogonal projection P_v. We further write $\Phi := \lambda\mathsf{Id} - A$ and consider the following commutative diagram:

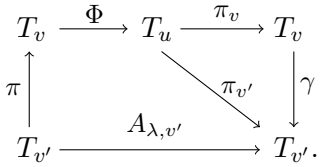

The top map $\pi_v \circ \Phi \colon T_v \to T_v$ equals $A_{\lambda,v}$ by definition. Our assumption $Av = \lambda v$ means $\Phi(v) = 0$. Hence $\Phi(x) = \Phi(\pi(x))$ for $x \in \mathbb{C}^n$. This implies that the bottom map $\pi_{v'} \circ \Phi \circ \pi$ indeed equals $A_{\lambda,v'}$. We conclude that

$$A_{\lambda,v'} = \gamma \circ A_{\lambda,v} \circ \pi, \quad A_{\lambda,v'}^{-1} = \pi^{-1} \circ A_{\lambda,v}^{-1} \circ \gamma^{-1}, \quad \|A_{\lambda,v'}^{-1}\| \le \|\pi^{-1}\| \cdot \|\gamma^{-1}\| \cdot \|A_{\lambda,v}^{-1}\|.$$

We will see in a moment that $\pi_{v'}^{-1}$ is bijective. Then, $\gamma^{-1} = \pi_v \circ \pi_{v'}^{-1}$, hence $\|\gamma^{-1}\| \le \|\pi_v\| \cdot \|\pi_{v'}^{-1}\|$.

We use the abbreviations $\delta := d_{\mathbb{P}}(v', v)$, $\delta_0 := d_{\mathbb{P}}(u, v)$, and $\delta_1 := d_{\mathbb{P}}(u, v')$. Note that $\delta_0 - \delta_1 \le \delta_1 \le \delta_0 + \delta$ by the triangle inequality. Let us proceed with some estimates.

Using the bound $\cos\delta \ge 1 - \frac{2}{\pi}\delta$, which is valid for $\delta \le \pi/2$, we get

$$\frac{\cos(\delta_0 + \delta)}{\cos\delta_0} = \cos\delta - \tan(\delta_0)\sin(\delta) \;\ge\; 1 - \Big(\frac{2}{\pi} + \tan\delta_0\Big)\delta. \tag{13}$$

This implies

$$\cos\delta \cdot \frac{\cos(\delta_0 + \delta)}{\cos\delta_0} \;\ge\; 1 - \Big(\frac{4}{\pi} + \tan\delta_0\Big)\delta.$$

We write $\mu := \mu(A, \lambda, v)$ to simplify the notation. Lemma 3.3 provides the following estimate

$$\tan\delta_0 \;\le\; \frac{1}{\cos\delta_0} = \sqrt{1 + \mu^2} \;\le\; \sqrt{3}\,\mu, \tag{14}$$

where the last inequality is due to Lemma 2.6. Again using Lemma 2.6, we estimate

$$\frac{4}{\pi} + \tan\delta_0 \;\le\; \Big(\frac{4\sqrt{2}}{\pi} + \sqrt{3}\Big)\,\mu \;\le\; 3.6\,\mu.$$

We conclude that

$$\cos\delta \cdot \frac{\cos(\delta_0 + \delta)}{\cos\delta_0} \;\ge\; 1 - 3.6\,\mu\delta > 0, \tag{15}$$

where the positivity is a consequence of our assumption. This shows that $\delta_1 \le \delta_0 + \delta < \pi/2$ and hence $\pi_{v'}$ is indeed bijective.

Lemma 4.2 yields

$$\|\pi_v\| \le \cos\delta_0, \quad \|\pi^{-1}\| \le \frac{1}{\cos\delta}, \quad \|\pi_{v'}^{-1}\| \le \frac{1}{\cos\delta_1} \le \frac{1}{\cos(\delta_0 + \delta)}.$$

Recall that $\|A^{-1}_{\lambda,v'}\| \le \|\pi^{-1}\| \cdot \|\pi_v\| \cdot \|\pi^{-1}_{v'}\| \cdot \|A^{-1}_{\lambda,v}\|$. We obtain

$$\|A^{-1}_{\lambda,v'}\| \;\le\; \frac{1}{\cos\delta} \cdot \frac{\cos\delta_0}{\cos(\delta_0+\delta)} \cdot \|A^{-1}_{\lambda,v}\|.$$

The estimate (15) yields

$$\|A^{-1}_{\lambda,v'}\| \;\le\; \frac{1}{1-3.6\mu\delta} \cdot \|A^{-1}_{\lambda,v}\|. \tag{16}$$

For the other inequality, we proceed similarly. We have

$$A_{\lambda,v} = \gamma^{-1} \circ A_{\lambda,v'} \circ \pi^{-1}, \quad A^{-1}_{\lambda,v} = \pi \circ A^{-1}_{\lambda,v'} \circ \gamma, \quad \|A^{-1}_{\lambda,v}\| \le \|\pi\| \cdot \|\gamma\| \cdot \|A^{-1}_{\lambda,v'}\|.$$

Moreover, $\|\gamma\| \le \|\pi_{v'}\| \cdot \|\pi_v^{-1}\|$. Lemma 4.2 yields

$$\|\pi_{v'}\| \le \cos\delta_1 \le \cos(\delta_0-\delta), \quad \|\pi_v^{-1}\| \le \frac{1}{\cos\delta_0}, \quad \|\pi\| \le \cos\delta.$$

This implies

$$\|A^{-1}_{\lambda,v}\| \;\le\; \cos\delta \cdot \frac{\cos(\delta_0-\delta)}{\cos\delta_0} \cdot \|A^{-1}_{\lambda,v'}\|.$$

Using (13) and (14) we estimate

$$\frac{\cos(\delta_0-\delta)}{\cos\delta_0} = \cos\delta' + \tan(\delta_0)\sin(\delta) \;\le\; 1+\sqrt{3}\,\mu\,\delta.$$

We arrive at

$$\|A^{-1}_{\lambda,v}\| \;\le\; (1+\sqrt{3}\,\mu\,\cdot\,\delta)\|A^{-1}_{\lambda,v'}\|,$$

which, together with (16) and $\|A\|_F = 1$, completes the proof. □

Proof of Proposition 4.1. The assumption

$$\mu(A,\lambda,v)\big(\|A'-A\|_F + |\lambda'-\lambda| + d_{\mathbb{P}}(v',v)\big) \;\le\; \frac{\varepsilon}{7.2} \tag{17}$$

yields $3.6\,\mu(A,\lambda,v)d_{\mathbb{P}}(v',v) \le \frac{\varepsilon}{2}$. Lemma 4.4 implies that

$$\mu(A,\lambda,v') \;\le\; \frac{\mu(A,\lambda,v)}{1-\varepsilon/2} \;\le\; 2\,\mu(A,\lambda,v), \tag{18}$$

where we have used $\varepsilon \le 1$ for the last equality. Combining this with (17), we conclude that

$$\mu(A,\lambda,v')\big(\|A'-A\|_F + |\lambda'-\lambda|\big) \;\le\; \frac{\varepsilon}{3.6}. \tag{19}$$

Lemma 4.3 and inequality (18) imply that

$$\mu(A',\lambda',v') \;\le\; \frac{\mu(A,\lambda,v')}{1-\varepsilon/3.6} \;\le\; \frac{\mu(A,\lambda,v)}{(1-\varepsilon/3.6)(1-\varepsilon/2)} \;\le\; (1+\varepsilon)\mu(A,\lambda,v),$$

the last since $\varepsilon \leq 0.37$. This proves one of the claimed inequalities.

For the other inequality we proceed similarly. We use again Lemma 4.4 now to obtain

$$\frac{\mu(A,\lambda,v)}{1+\varepsilon/3.6} \;\leq\; \mu(A,\lambda,v'). \tag{20}$$

Inequality (19) allows us to use again Lemma 4.3, now to obtain

$$(1-\varepsilon/2)\,\mu(A,\lambda,v') \;\leq\; \mu(A',\lambda',v'). \tag{21}$$

Combining (20) and (21) we obtain

$$\mu(A',\lambda',v') \geq (1-\varepsilon/2)\,\mu(A,\lambda,v') \geq \frac{1-\varepsilon/2}{1+\varepsilon/3.6}\,\mu(A,\lambda,v) \geq \frac{1}{1+\varepsilon}\,\mu(A,\lambda,v)$$

the last inequality, again, since $\varepsilon \leq 0.37$. This completes the proof. □

5 Proof of Proposition 2.19

The proof relies on three main ingredients: Proposition 2.7, and Theorems 2.9 and 2.14. Different versions of it have been written down for different contexts: complex polynomial systems ([6, Theorem 3.1] or [7, Theorem 17.3]), ditto but with finite precision ([5, Theorem 4.3]), and Hermitian matrices ([2, Proposition 4]).

We set $\varepsilon := 0.12$ and $C_\varepsilon := \frac{\varepsilon}{12.5} = 0.0096$. Furthermore, we let $\xi := \frac{2C_\varepsilon(1-\varepsilon)}{3\sqrt{6}(1+\varepsilon)^4} \approx 0.001461$. Also, as in §2.8, we write $\alpha = d_{\mathbb{S}}(M,A)$.

Let $0 = \tau_0 < \tau_1 < \cdots < \tau_K = 1$ and $(\lambda_0, v_0) = (\zeta_0, w_0), (\zeta_1, w_1), \ldots, (\zeta_K, w_K)$ be the sequences of τ-values and pairs in $\mathbb{C} \times \mathbb{C}^n$ generated by the algorithm Path-follow. We simplify notation and write Q_i instead of Q_{τ_i} and (λ_i, v_i) instead of $(\lambda_{\tau_i}, v_{\tau_i})$. (There is no danger of confusing this (λ_1, v_1) with the one appearing in the statement of Proposition 2.19.) We associate with the solution path (11) in $\mathcal{V}$ the following curve in $\mathbb{S}(\mathbb{C}^{n\times n}) \times \mathbb{C} \times \mathbb{C}^n$:

$$[0,1] \to V, \quad \tau \mapsto (P_\tau, \widehat{\lambda_\tau}, v_\tau) := \Big(\frac{Q_\tau}{\|Q_\tau\|_F}, \frac{\lambda_\tau}{\|Q_\tau\|_F}, v_\tau\Big). \tag{22}$$

We also write P_i instead of P_{τ_i}. The meaning of the parameterization by τ is that $\alpha\tau$ is the parameterization of $\tau \mapsto P_\tau$ by arc length, which means that $\left\|\frac{dP_\tau}{d\tau}\right\| = \alpha$.

We will carry out the proof on the curve (22) in the sphere $\mathbb{S}(\mathbb{C}^{n\times n})$. We do so to simplify the exposition and without implying that algorithm Path-follow should be modified to normalize matrices. Indeed, all the functions on triples in $\mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{C}^n$ involved in our proof —dist, μ, and Newton's operator— are scale invariant on the first two components. Furthermore, to avoid burdening the notation, we will write λ instead of $\widehat{\lambda}$. This should introduce no confusion.

The following result is the technical core in the proof of Proposition 2.19.

Proposition 5.1 *For $i = 0, \ldots, K-1$, the following statements are true:*

(a) $\mathsf{dist}((P_i, \zeta_i, w_i), (P_i, \lambda_i, v_i)) \leq \dfrac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)}$.

(b) $\dfrac{\mu(P_i, \zeta_i, w_i)}{1+\varepsilon} \leq \mu(P_i, \lambda_i, v_i) \leq (1+\varepsilon)\mu(P_i, \zeta_i, w_i)$.

(c) $\mathsf{dist}((P_i, \lambda_i, v_i), (P_{i+1}, \lambda_{i+1}, v_{i+1})) \leq \dfrac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} \dfrac{2(1-\varepsilon)}{3(1+\varepsilon)}$.

(d) $\mathsf{dist}((P_{i+1}, \zeta_i, w_i), (P_{i+1}, \lambda_{i+1}, v_{i+1})) \leq \dfrac{2C_\varepsilon}{(1+\varepsilon)\mu(P_i, \lambda_i, v_i)}$.

(e) *(ζ_i, w_i) is an approximate eigenpair of P_{i+1} with associated eigenpair (λ_{i+1}, v_{i+1}).* □

PROOF. We proceed by induction, showing that

$$(\mathbf{a}, i) \Rightarrow (\mathbf{b}, i) \Rightarrow (\mathbf{c}, i) \Rightarrow (\mathbf{d}, i) \Rightarrow \big((\mathbf{e}, i) \text{ and } (\mathbf{a}, i+1)\big).$$

Inequality (a) for $i = 0$ is trivial.

Assume now that (a) holds for some $i \leq K-1$. Then, Theorem 2.9 (with $A = A' = P_i$) implies

$$\frac{\mu(P_i, \zeta_i, w_i)}{1+\varepsilon} \leq \mu(P_i, \lambda_i, v_i) \leq (1+\varepsilon)\mu(P_i, \zeta_i, w_i)$$

and thus (b). We now prove (c). To do so, let $\tau_* > \tau_i$ be such that

$$\int_{\tau_i}^{\tau_*} \left\| \frac{d(P_\tau, \lambda_\tau, v_\tau)}{d\tau} \right\| d\tau = \frac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} \frac{2(1-\varepsilon)}{3(1+\varepsilon)}$$

or $\tau_* = 1$, whichever is smaller. Then, for all $t \in [\tau_i, \tau_*]$,

$$\begin{aligned} \mathsf{dist}((P_i, \lambda_i, v_i), (P_t, \lambda_t, v_t)) &\leq d_{\mathsf{R}}((P_i, \lambda_i, v_i), (P_t, \lambda_t, v_t)) \\ &= \int_{\tau_i}^{t} \left\| \frac{d(P_\tau, \lambda_\tau, v_\tau)}{d\tau} \right\| d\tau \leq \frac{C_\varepsilon}{\mu(P_i, \lambda_i, v_i)} \frac{2(1-\varepsilon)}{3(1+\varepsilon)}, \end{aligned} \tag{23}$$

the first inequality by (2). It is therefore enough to show that $\tau_{i+1} \leq \tau_*$. This is trivial if $\tau_* = 1$. We therefore assume $\tau_* < 1$. The bound above allows us to apply Theorem 2.9 and to deduce, for all $\tau \in [\tau_i, \tau_*]$,

$$\frac{1}{1+\varepsilon}\, \mu(P_i, \lambda_i, v_i) \leq \mu(P_\tau, \lambda_\tau, v_\tau) \leq (1+\varepsilon)\, \mu(P_i, \lambda_i, v_i). \tag{24}$$

Proposition 2.7 implies that

$$\left\| \frac{d}{d\tau}(P_\tau, \lambda_\tau, v_\tau) \right\| \leq \sqrt{6}\, \mu(P_\tau, \lambda_\tau, v_\tau) \left\| \frac{d}{d\tau} P_\tau \right\|$$

We now use (24) to deduce that

$$\frac{C_\varepsilon}{\mu(P_i,\lambda_i,v_i)}\frac{2(1-\varepsilon)}{3(1+\varepsilon)} \;=\; \int_{\tau_i}^{\tau_*}\left\|\frac{d(P_\tau,\lambda_\tau,v_\tau)}{d\tau}\right\|d\tau \;\le\; \int_{\tau_i}^{\tau_*}\sqrt{6}\,\mu(P_\tau,\lambda_\tau,v_\tau)\left\|\frac{d}{d\tau}P_\tau\right\|d\tau$$
$$\le\; \sqrt{6}\,(1+\varepsilon)\,\mu(P_i,\lambda_i,v_i)\int_{\tau_i}^{\tau_*}\left\|\frac{d}{d\tau}P_\tau\right\|d\tau \;=\; \sqrt{6}\,(1+\varepsilon)\,\mu(P_i,\lambda_i,v_i)\,d_{\mathbb{S}}(P_i,P_{\tau_*}).$$

Consequently, using (b), we obtain

$$d_{\mathbb{S}}(P_i,P_{\tau_*}) \ge \frac{2C_\varepsilon(1-\varepsilon)}{3\sqrt{6}(1+\varepsilon)^2\mu^2(P_i,\lambda_i,v_i)} \ge \frac{2C_\varepsilon(1-\varepsilon)}{3\sqrt{6}(1+\varepsilon)^4\mu^2(P_i,\zeta_i,w_i)}.$$

Recall that the parameter ξ in Path-follow was chosen as $\xi = \frac{2C_\varepsilon(1-\varepsilon)}{3\sqrt{6}(1+\varepsilon)^4}$. By the definition of $\tau_{i+1}-\tau_i$ in Path-follow we have $\alpha(\tau_{i+1}-\tau_i) = \frac{\xi}{\mu^2(P_i,\zeta_i,w_i)}$. So we obtain

$$d_{\mathbb{S}}(P_i,P_{\tau_*}) \ge \alpha(\tau_{i+1}-\tau_i) \;=\; d_{\mathbb{S}}(P_i,P_{i+1}).$$

This implies $\tau_{i+1} \le \tau_*$ as claimed, and hence inequality (c) follows from (23) with $t = \tau_{i+1}$. With it, we may apply Theorem 2.9 once more to deduce, for all $\tau \in [\tau_i,\tau_{i+1}]$,

$$\frac{\mu(P_i,\lambda_i,v_i)}{1+\varepsilon} \le \mu(P_\tau,\lambda_\tau,v_\tau) \le (1+\varepsilon)\mu(P_i,\lambda_i,v_i). \tag{25}$$

We now observe that

$$\begin{aligned}\mathsf{dist}((P_{i+1},\zeta_i,w_i),(P_i,\zeta_i,w_i)) &= \|P_{i+1}-P_i\|_F \;\le\; d_{\mathbb{S}}(P_i,P_{i+1}) \;=\; \alpha(\tau_{i+1}-\tau_i)\\ &= \frac{\xi}{\mu^2(P_i,\zeta_i,w_i)} \;\le\; \frac{2C_\varepsilon(1-\varepsilon)}{3\sqrt{6}(1+\varepsilon)\mu(P_i,\zeta_i,w_i)}\end{aligned}$$

and use this bound, together with the triangle inequality, (24), (a), and (c) to obtain

$$\begin{aligned}\mathsf{dist}((P_{i+1},&\zeta_i,w_i),(P_{i+1},\lambda_{i+1},v_{i+1})) \le \mathsf{dist}((P_{i+1},\zeta_i,w_i),(P_i,\zeta_i,w_i))\\ &+\mathsf{dist}((P_i,\zeta_i,w_i),(P_i,\lambda_i,v_i))\\ &+\mathsf{dist}((P_i,\lambda_i,v_i),(P_{i+1},\lambda_{i+1},v_{i+1}))\\ &\le \frac{2C_\varepsilon(1-\varepsilon)}{3\sqrt{6}(1+\varepsilon)\mu(P_i,\zeta_i,w_i)} + \frac{C_\varepsilon}{\mu(P_i,\lambda_i,v_i)} + \frac{C_\varepsilon}{\mu(P_i,\lambda_i,v_i)}\frac{2}{3}\frac{1-\varepsilon}{1+\varepsilon}\\ &\le \frac{2C_\varepsilon}{(1+\varepsilon)\mu(P_i,\lambda_i,v_i)}\end{aligned} \tag{26}$$

which proves (d). We now note that $\frac{2C_\varepsilon}{1+\varepsilon} < c_0 = 0.2881$. We can therefore apply Theorem 2.14 to deduce that (ζ_i,w_i) is an approximate eigenpair of P_{i+1} associated with its eigenpair (λ_{i+1},v_{i+1}), and hence (e) holds.

It follows from (e) that $(\zeta_{i+1},w_{i+1}) = N_{P_{i+1}}(\zeta_i,w_i)$ satisfies

$$\mathsf{dist}((P_{i+1},\zeta_{i+1},w_{i+1}),(P_{i+1},\lambda_{i+1},v_{i+1})) \le \frac{1}{2}\,\mathsf{dist}((P_{i+1},\zeta_i,w_i),(P_{i+1},\lambda_{i+1},v_{i+1})).$$

Using this bound, (d) and the right-hand inequality in (25) with $\tau = \tau_{i+1}$, we obtain

$$\mathsf{dist}((P_{i+1}, \zeta_{i+1}, w_{i+1}), (P_{i+1}\lambda_{i+1}, v_{i+1})) \leq \frac{C_\varepsilon}{(1+\varepsilon)\mu(P_i, \lambda_i, v_i)} \leq \frac{C_\varepsilon}{\mu(P_{i+1}, \lambda_{i+1}, v_{i+1})},$$

which proves (a) for $i+1$. The proposition is thus proved. □

The following is an immediate consequence of Proposition 5.1(c) and Theorem 2.9.

Corollary 5.2 *For all $i = 0, \ldots, K-1$ and $\tau \in [\tau_i, \tau_{i+1}]$, we have*

$$\frac{\mu(Q_i, \lambda_i, v_i)}{1+\varepsilon} \leq \mu(Q_\tau, \lambda_\tau, v_\tau) \leq (1+\varepsilon)\mu(Q_i, \lambda_i, v_i). \qquad \square$$

Proof of Proposition 2.19. It follows from Proposition 5.1(e) for $i = K-1$ that (ζ_{K-1}, w_{K-1}) is an approximate eigenpair of $Q_K = A$ with associated eigenpair (λ_K, v_K). Consequently, so is the returned point $(\zeta_K, w_K) = N_A(\zeta_{K-1}, w_{K-1})$.

Consider now any $i \in \{0, \ldots, K-1\}$. Using Corollary 5.2, and Proposition 5.1(b), and by the choice of the step size $\Delta\tau$ in Algorithm 1, we obtain

$$\begin{aligned}
\int_{\tau_i}^{\tau_{i+1}} \mu^2(Q_\tau, \lambda_\tau, v_\tau) d\tau &\geq \int_{\tau_i}^{\tau_{i+1}} \frac{\mu^2(Q_i, \lambda_i, v_i)}{(1+\varepsilon)^2} d\tau = \frac{\mu^2(Q_i, \lambda_i, v_i)}{(1+\varepsilon)^2}(\tau_{i+1} - \tau_i) \\
&\geq \frac{\mu^2(Q_i, \zeta_i, w_i)}{(1+\varepsilon)^4}(\tau_{i+1} - \tau_i) \\
&= \frac{\mu^2(Q_i, \zeta_i, w_i)}{(1+\varepsilon)^4} \frac{\xi}{\alpha\mu^2(Q_i, \zeta_i, w_i)} = \frac{\xi}{(1+\varepsilon)^4\alpha} \\
&\geq \frac{1}{1077\,\alpha}.
\end{aligned}$$

This implies

$$\int_0^1 \mu^2(Q_\tau, \lambda_\tau, v_\tau) d\tau \geq \frac{K}{1077\,\alpha},$$

which proves the stated upper bound on K. The lower bound follows from

$$\begin{aligned}
\int_{\tau_i}^{\tau_{i+1}} \mu^2(Q_\tau, \lambda_\tau, v_\tau) d\tau &\leq \int_{\tau_i}^{\tau_{i+1}} \mu^2(Q_i, \lambda_i, v_i)(1+\varepsilon)^2 d\tau \\
&= \mu^2(Q_i, \lambda_i, v_i)(1+\varepsilon)^2(\tau_{i+1} - \tau_i) \\
&\leq \mu^2(Q_i, \zeta_i, w_i)(1+\varepsilon)^4(\tau_{i+1} - \tau_i) \\
&= \frac{\xi(1+\varepsilon)^4}{\alpha} \leq \frac{1}{434\,\alpha}.
\end{aligned}$$

□

As pointed out at the beginning of this section, the proof above follows the steps of the one in [2, Proposition 4]. In both cases we obtain

$$\frac{3}{C_\varepsilon(1-\varepsilon)} \leq \frac{K}{\alpha \int_0^1 \mu^2(Q_\tau, \lambda_\tau, v_\tau) d\tau} \leq \frac{3(1+\varepsilon)^8}{C_\varepsilon(1-\varepsilon)}.$$

The difference in the actual constants in both statements is due to the difference between the values of ε and C_ε in (the inadecuate) [1, Proposition 3.22] used in [2] and its improved version, Theorem 2.9 (see the paragraph following Proposition 3 in [2]).

6 Proof of Theorem 2.16

This section is the technical heart of the paper. We divide it into several subsections, the first of which summarizes some notions and tools of probability theory on Riemannian manifolds.

6.1 The coarea formula

On a Riemannian manifold M there is a well-defined measure vol_M obtained by integrating the indicator functions $\mathbb{1}_A$ of Borel-measurable subsets $A \subseteq M$ against the volume form dM of M:

$$\mathrm{vol}_M(A) = \int_M \mathbb{1}_A \, dM.$$

Dividing $\mathbb{1}$ by $\mathrm{vol}_M(M)$ if $\mathrm{vol}_M(M) < \infty$, this leads to a natural notion of *uniform distribution* on M, which we will denote by $\mathscr{U}(M)$. More generally, we will call any measurable function $f\colon M \to [0,\infty]$ such that $\int_M f \, dM = 1$ a *probability density* on M.

The coarea formula is an extension of the transformation formula to not necessarily bijective smooth maps between Riemannian manifolds. In order to state it, we first need to generalize the notion of Jacobians.

Suppose that M, N are Riemannian manifolds of dimensions m, n, respectively such that $m \geq n$. Let $\psi\colon M \to N$ be a smooth map. By definition, the derivative $D\psi(x)\colon T_xM \to T_{\psi(x)}N$ at a regular point $x \in M$ is surjective. Hence the restriction of $D\psi(x)$ to the orthogonal complement of its kernel yields a linear isomorphism. The absolute value of its determinant is called the *normal Jacobian* of ψ at x and denoted by $\mathrm{NJ}\psi(x)$. We set $\mathrm{NJ}\psi(x) := 0$ if x is not a regular point.

If y is a regular value of ψ, then the fiber $F_y := \psi^{-1}(y)$ is a Riemannian submanifold of M of dimension $m - n$. Sard's lemma states that almost all $y \in N$ are regular values.

We can now state the *coarea formula*.

Theorem 6.1 (Coarea formula) *Suppose that M, N are Riemannian manifolds of dimensions m, n, respectively, and let $\psi\colon M \to N$ be a surjective smooth map. Put $F_y = \psi^{-1}(y)$. Then we have for any function $\chi\colon M \to \mathbb{R}$ that is integrable with respect to the volume measure of M that*

$$\int_M \chi \, dM = \int_{y \in N} \left(\int_{F_y} \frac{\chi}{\mathrm{NJ}\psi} \, dF_y \right) dN. \qquad \square$$

It should be clear that this result contains the transformation formula as a special case. Moreover, if we apply the coarea formula to the projection $\pi_2 \colon M \times N \to N$, $(x, y) \mapsto y$, we retrieve Fubini's equality since $\mathrm{NJ}\pi_2 = 1$.

The coarea formula is useful to define the concepts of marginal and conditional distributions for densities defined on a product space $M \times N$ when the components are Riemannian manifolds.

Suppose that we are in the situation described in the statement of Theorem 6.1 and we have a probability measure on M with density ρ_M. For a regular value $y \in N$ we set

$$\rho_N(y) := \int_{F_y} \frac{\rho_M}{\mathrm{NJ}\psi} \, dF_y. \tag{27}$$

The coarea formula implies that for all measurable sets $B \subseteq N$ we have

$$\int_{\psi^{-1}(B)} \rho_M \, dM = \int_B \rho_N \, dN.$$

Hence ρ_N is a probability density on N. We call it the *pushforward* of ρ_M with respect to ψ. In the special case that $\psi \colon M \times N \to N$, $(x, y) \mapsto y$, is the projection, we have $\mathrm{NJ}\psi = 1$, and we retrieve the usual formula for the marginal density.

Furthermore, for a regular value $y \in N$ and $x \in F_y$ we define the *conditional density* on F_y

$$\rho_{F_y}(x) := \frac{\rho_M(x)}{\rho_N(y)\mathrm{NJ}\psi(x)}. \tag{28}$$

Clearly, this defines a probability density on F_y. Again, in the special case that $\psi \colon M \times N \to N$, $(x, y) \mapsto y$, we retrieve the usual formula for the conditional density.

The coarea formula implies that for all measurable functions $\chi \colon M \to \mathbb{R}$,

$$\int_M \chi \, \rho_M \, dM = \int_{y \in N} \left(\int_{F_y} \chi \, \rho_{F_y} \, dF_y \right) \rho_N(y) \, dN,$$

provided the left-hand integral exists. Therefore, we can interpret ρ_{F_y} as the *density of the conditional distribution* of x on the fiber F_y and briefly express the formula above in probabilistic terms as

$$\mathop{\mathbb{E}}_{x \sim \rho_M} \chi(x) = \mathop{\mathbb{E}}_{y \sim \rho_N} \mathop{\mathbb{E}}_{x \sim \rho_{F_y}} \chi(x). \tag{29}$$

6.2 An auxiliary result from linear algebra

Let E and F be finite dimensional Euclidean vector spaces such that $\dim E \geq \dim F$. If $\varphi \colon E \to F$ is a surjective linear map, we denote by $\widetilde{\varphi} \colon (\ker \varphi)^\perp \to F$ the restriction

of φ to the orthogonal complement of the kernel of φ. Then $\widetilde{\varphi}$ is surjective and we define its *normal determinant* by

$$\operatorname{ndet}\varphi := |\det\widetilde{\varphi}|.$$

(If φ is not surjective, we set $\operatorname{ndet}\varphi := 0$.) We consider the graph $\Gamma := \{(x,\varphi(x)) \mid x \in E\}$ of φ. Then, Γ is a linear subspace of $E \times F$ and the two projections

$$p_1\colon \Gamma \to E,\ (x,\varphi(x)) \mapsto x, \quad p_2\colon \Gamma \to F,\ (x,\varphi(x)) \mapsto \varphi(x)$$

are linear maps. Note that p_1 is an isomorphism and p_2 is surjective as φ is so.

Lemma 6.2 *Under the above assumptions, we have*

$$\frac{\operatorname{ndet} p_1}{\operatorname{ndet} p_2} = (\operatorname{ndet}\varphi)^{-1}.$$

Proof. Let $K := \ker\varphi$ and $\widetilde{E}$ be the orthogonal complement of K in E. Let $\widetilde{\Gamma} \subseteq \widetilde{E} \times F$ denote the graph of $\widetilde{\varphi}\colon \widetilde{E} \to F$. Further, let $\widetilde{p}_1\colon \widetilde{\Gamma} \to \widetilde{E}$ and $\widetilde{p}_2\colon \widetilde{\Gamma} \to F$ denote the projections. Since φ is surjective, $\widetilde{\varphi}$ is bijective. We have obvious isometries $\Gamma \simeq \widetilde{\Gamma} \times K$, $p_1 \simeq \widetilde{p}_1 \times \mathsf{Id}_K$, and we can interpret $\widetilde{\Gamma}$ as the orthogonal complement of $\ker p_2 = K \times \{0\}$ in Γ. By the definition of the normal determinant, we have

$$\operatorname{ndet}\varphi = |\det\widetilde{\varphi}|, \quad \operatorname{ndet} p_1 = |\det\widetilde{p}_1|, \quad \operatorname{ndet} p_2 = |\det\widetilde{p}_2|.$$

It is therefore sufficient to prove that

$$\Big|\frac{\det\widetilde{p}_1}{\det\widetilde{p}_2}\Big| = |\det\widetilde{\varphi}|^{-1}.$$

The singular value decomposition tells us that, with respect to suitable orthonormal bases on $\widetilde{E}$ and F, the representation matrix of $\widetilde{\varphi}$ equals $\mathsf{diag}(\sigma_1,\dots,\sigma_n)$, where $\sigma_1 \geq \cdots \geq \sigma_n$ are the singular values of $\widetilde{\varphi}$. Note that $|\det\widetilde{\varphi}| = \sigma_1 \cdots \sigma_n$. It is now straightforward to check that

$$|\det\widetilde{p}_1| = \prod_{i=1}^{n} \frac{1}{\sqrt{1+\sigma_i^2}}, \quad |\det\widetilde{p}_2| = \prod_{i=1}^{n} \frac{\sigma_i}{\sqrt{1+\sigma_i^2}}.$$

Therefore,

$$\Big|\frac{\det\widetilde{p}_1}{\det\widetilde{p}_2}\Big| = \prod_{i=1}^{n} \sigma_i^{-1} = |\det\widetilde{\varphi}|^{-1},$$

which finishes the proof. □

Suppose that W is a finite dimensional complex vector space with a Hermitian inner product $\langle\ ,\ \rangle_{\mathbb{C}}$. The real part of $\langle\ ,\ \rangle_{\mathbb{C}}$ turns W into a Euclidean vector space. Let $\psi\colon W \to W$ be a $\mathbb{C}$-linear map. If we denote by $\det\psi_{\mathbb{R}}$ the determinant of ψ, considered as an $\mathbb{R}$-linear map, then it is a well known fact that $\det\psi_{\mathbb{R}} = |\det\psi|^2$. (Indication of proof: the singular value decomposition allows to reduce to the case $W = \mathbb{C}$ and $\psi(z) = \lambda z$. Since ψ is the composition of a rotation and a homothety by the stretching factor $|\lambda|$, it follows that $\det\psi_{\mathbb{R}} = |\lambda|^2$.)

6.3 Normal Jacobians for the eigenpair problem

Recall from §2.1, we have the two projections

$$\pi_1 \colon \mathcal{V} \to X, (A, \lambda, [v]) \to A, \quad \pi_2 \colon \mathcal{V} \to Y, (A, \lambda, [v]) \to (\lambda, [v])$$

and, for $(A, \lambda, v) \in \mathcal{V}$, the linear operator $A_{\lambda,v} \colon T_v \to T_v$ given by $P_{v^\perp}(A - \lambda \,\mathsf{Id})_{|T_v}$.

Proposition 6.3 *Let $p := (A, \lambda, [v]) \in \mathcal{V}$. Then λ is a simple eigenvalue of A iff $A_{\lambda,v}$ is invertible. In this case, the derivative $D\pi_1(p) \colon T_p \to T_A X$ is an isomorphism, the derivative $D\pi_2(p) \colon T_p \to T_{(\lambda,[v])}Y$ is surjective, and we have*

$$\frac{\mathrm{NJ}\pi_1(p)}{\mathrm{NJ}\pi_2(p)} = |\det A_{\lambda,v}|^2 = \det(A_{\lambda,v} A^*_{\lambda,v}).$$

PROOF. Let $p := (A, \lambda, [v]) \in \mathcal{V}$. We suppose that $\|v\| = 1$ and we identify the tangent space $T_{[v]}\mathbb{P}(\mathbb{C}^n)$ with T_v. By orthogonal invariance, we may assume without loss of generality that $v = e_1 = (1, 0, \ldots, 0)$. Then we write

$$A = \begin{bmatrix} \lambda & c^T \\ 0 & B \end{bmatrix}, \quad c \in \mathbb{C}^{n-1}, B \in \mathbb{C}^{(n-1)\times(n-1)}.$$

The matrix $M := \lambda I - B$ represents the linear map $A_{\lambda,v}$. Clearly, M is invertible iff λ is a simple eigenvalue of A, which shows the first assertion of the proposition. We assume now that M is invertible.

Let $(\dot{A}, \dot{\lambda}, \dot{v}) \in \mathbb{C}^{n\times n} \times \mathbb{C} \times T_v$. According to equation (14.19) in [7], the tangent space of $\mathcal{V}$ at p is characterized in the following way:

$$(\dot{A}, \dot{\lambda}, \dot{v}) \in T_p\mathcal{V} \iff \dot{A}v + (A - \lambda I)\dot{v} - \dot{\lambda}v = 0. \tag{30}$$

In order to express $\dot{\lambda}$, $\dot{v}$ in terms of $\dot{A}$, we denote by $\dot{a}_i \in \mathbb{C}^n$ the ith column of A, and we write $\dot{a}_1 = (\dot{a}, \dot{b})$ where $\dot{a} \in \mathbb{C}$ and $\dot{b} \in \mathbb{C}^{n-1}$. Also, since $\dot{v} \in T_v$ (and we are assuming $v = e_1$) we have $\dot{v} = (0, \dot{w})$ for some $\dot{w} \in \mathbb{C}^{n-1}$. Using this notation, equation (30) can be rewritten as

$$\begin{aligned} \dot{a} + c^T\dot{w} - \dot{\lambda} &= 0 \\ \dot{b} - M\dot{w} &= 0. \end{aligned}$$

This system of equations has the unique solution

$$\dot{w} = M^{-1}\dot{b}, \quad \dot{\lambda} = \dot{a} + c^T M^{-1}\dot{b}. \tag{31}$$

So we can interpret $T_p\mathcal{V}$ as the graph Γ of the linear map $\varphi \colon E \to F$ (with $E = \mathbb{C}^{n\times n}$ and $F = \mathbb{C} \times T_v$) given by $\dot{A} \mapsto (\dot{\lambda}, \dot{v})$, and $D\pi_1(p), D\pi_2(p)$ are the corresponding projections $p_1 \colon \Gamma \to E$ and $p_2 \colon \Gamma \to F$, respectively. According to Lemma 6.2, it therefore suffices to prove that

$$\mathrm{ndet}\,\varphi = |\det A_{\lambda,v}|^{-2}. \tag{32}$$

Since $\dot{w} = 0$, $\dot{\lambda} = 0$ implies $\dot{b} = 0$, $\dot{a} = 0$, we see that the orthogonal complement of the kernel of φ is given by the conditions $\dot{a}_2 = 0, \ldots, \dot{a}_n = 0$. The restriction $\widetilde{\varphi}$ of φ to $(\ker \varphi)^\perp$, $(\dot{a}, \dot{b}) \mapsto (\dot{\lambda}, \dot{w})$, according to (31), has the following matrix

$$\begin{bmatrix} 1 & c^T M^{-1} \\ 0 & M^{-1} \end{bmatrix}$$

with respect to the standard bases. Therefore, $\det \widetilde{\varphi} = \det M^{-1}$. The determinant of $\widetilde{\varphi}$, seen as a $\mathbb{R}$-bilinear map, therefore equals $|\det M^{-1}|^2$, see the comment at the end of §6.2. We conclude that

$$\operatorname{ndet} \varphi = |\det M|^{-2} = |\det A_{\lambda,v}|^{-2},$$

which proves the claimed equality (32). □

6.4 Orthogonal decompositions

We shall distinguish points in $\mathbb{P}(\mathbb{C}^n)$ from their representatives in $\mathbb{S}(\mathbb{C}^n)$ and, accordingly, consider the following lifting of the solution variety $\mathcal{V}$

$$\hat{\mathcal{V}} := \{(A, \lambda, v) \in \mathbb{C}^{n\times n} \times \mathbb{C} \times \mathbb{S}(\mathbb{C}^n) \mid Av = \lambda v\}.$$

Abusing notation, we denote the projection $\hat{\mathcal{V}} \to \mathbb{C} \times \mathbb{S}(\mathbb{C}^n)$ by π_2 as well. The fiber of π_2 at $(\lambda, v) \in \mathbb{C} \times \mathbb{S}(\mathbb{C}^n)$ equals

$$V_{(\lambda,v)} = \{A \in \mathbb{C}^{n\times n} \mid Av = \lambda v\}.$$

This is an affine linear subspace of $\mathbb{C}^{n\times n}$ with the corresponding linear space $V^{\mathsf{lin}}_{(\lambda,v)} = \{A \in \mathbb{C}^{n\times n} \mid Av = 0\}$. We denote by $C_{(\lambda,v)}$ the orthogonal complement of $V^{\mathsf{lin}}_{(\lambda,v)}$ in $\mathbb{C}^{n\times n}$. So we have the orthogonal decomposition

$$\mathbb{C}^{n\times n} = V^{\mathsf{lin}}_{(\lambda,v)} \oplus C_{(\lambda,v)}. \tag{33}$$

Let $K_{(\lambda,v)}$ denote the point in $V_{(\lambda,v)}$ that is closest to the origin (with respect to the Frobenius norm). Note that $K_{(\lambda,v)} \in C_{(\lambda,v)}$ and $V_{(\lambda,v)} = V^{\mathsf{lin}}_{(\lambda,v)} + K_{(\lambda,v)}$.

Recall, $\varphi^{\overline{A},\sigma}_{n\times n}$ denotes the density of the Gaussian $\mathcal{N}(\overline{A}, \sigma^2\mathsf{Id})$ on $\mathbb{C}^{n\times n}$, where $\overline{A} \in \mathbb{C}^{n\times n}$ and $\sigma > 0$. For fixed $(\lambda, v) \in \mathbb{C} \times \mathbb{S}(\mathbb{C}^n)$, we decompose the mean $\overline{A}$ acccording to (33) as

$$\overline{A} = \overline{M}_{(\lambda,v)} + \overline{K}_{(\lambda,v)}$$

where $\overline{M}_{(\lambda,v)} \in V^{\mathsf{lin}}_{(\lambda,v)}$ and $\overline{K}_{(\lambda,v)} \in C_{(\lambda,v)}$. If we denote by $\varphi_{V^{\mathsf{lin}}_{(\lambda,v)}}$ and $\varphi_{C_{(\lambda,v)}}$ the densities of the Gaussian distributions in the spaces $V^{\mathsf{lin}}_{(\lambda,v)}$ and $C_{(\lambda,v)}$ with covariance matrices $\sigma^2\mathsf{Id}$ and means $\overline{M}_{(\lambda,v)}$ and $\overline{K}_{(\lambda,v)}$, respectively, then the density $\varphi^{\overline{A},\sigma}_{n\times n}$ factors as

$$\varphi^{\overline{A},\sigma}_{n\times n}(M + K) = \varphi_{V^{\mathsf{lin}}_{(\lambda,v)}}(M) \cdot \varphi_{C_{(\lambda,v)}}(K). \tag{34}$$

Consider now the projection $\pi_1\colon \hat{\mathcal{V}} \to \mathbb{C}^{n\times n}$. Its fiber V_A at $A \in \mathbb{C}^{n\times n} \setminus \Sigma$ is a disjoint union of n unit circles and therefore has volume $2\pi n$. We associate with $\varphi_{n\times n}^{\overline{A},\sigma}$ the function $\rho_{\hat{\mathcal{V}}} : \hat{\mathcal{V}} \to \mathbb{R}$ defined by

$$\rho_{\hat{\mathcal{V}}}(A,\lambda,v) := \frac{1}{2\pi n}\, \varphi_{n\times n}^{\overline{A},\sigma}(A)\, \mathrm{NJ}\pi_1(A,\lambda,v). \tag{35}$$

The proof of the following result is done as in [7, Lemma 18.10].

Lemma 6.4 (a) *The function $\rho_{\hat{\mathcal{V}}}$ is a probability density on $\hat{\mathcal{V}}$.*

(b) *The expectation of a function $F\colon \hat{\mathcal{V}} \to \mathbb{R}$ with respect to $\rho_{\hat{\mathcal{V}}}$ can be expressed as*

$$\mathop{\mathbb{E}}_{(A,\lambda,v)\sim\rho_{\hat{\mathcal{V}}}} F(A,\lambda,v) = \mathop{\mathbb{E}}_{A\sim\varphi_{n\times n}^{\overline{Q},\sigma}} F_{\mathsf{sav}}(A),$$

where $F_{\mathsf{sav}}(A) := \frac{1}{2\pi n}\int_{V_A} F\, dV_A$.

(c) *The pushforward of $\rho_{\hat{\mathcal{V}}}$ with respect to $\pi_1\colon \hat{\mathcal{V}} \to \mathbb{C}^{n\times n}$ equals $\varphi_{n\times n}^{\overline{A},\sigma}$.*

(d) *For $A \notin \Sigma$, the conditional density on the fiber V_A is the density of the uniform distribution on V_A.* □

Remark 6.5 In the particular case that $F\colon \hat{\mathcal{V}} \to \mathbb{R}$ is given by $F(A,\lambda,v) = \frac{\mu^2(A,\lambda,v)}{\|A\|_F^2}$ we have $F_{\mathsf{sav}}(A) = \frac{\mu_{\mathsf{av}}^2(A)}{\|A\|_F^2}$.

Recall from (4) that we assigned to $A \in \mathbb{C}^{n\times n}$ the linear operator $A_{\lambda,v}\colon T_v \to T_v$ obtained by restricting $A - \lambda\mathsf{Id}$ to T_v. We write

$$c_{(\lambda,v)} := \int_{M\in V_{(\lambda,v)}^{\mathsf{lin}}} |\det M_{\lambda,v}|^2\, \varphi_{V_{(\lambda,v)}^{\mathsf{lin}}}(M)\, dM. \tag{36}$$

Lemma 6.6 *The pushforward density $\rho_{\mathbb{C}\times\mathbb{S}(\mathbb{C}^n)}$ of $\rho_{\hat{\mathcal{V}}}$ with respect to $\pi_2\colon \hat{\mathcal{V}} \to \mathbb{C}\times\mathbb{S}(\mathbb{C}^n)$ is given by*

$$\rho_{\mathbb{C}\times\mathbb{S}(\mathbb{C}^n)}(\lambda,v) = \frac{c_{(\lambda,v)}}{2\pi n} \cdot \varphi_{C_{(\lambda,v)}}(K_{(\lambda,v)}), \tag{37}$$

the conditional density $\widetilde{\rho}_{V_{(\lambda,v)}}$ on the fiber $V_{(\lambda,v)}$ of π_2 is given by

$$\widetilde{\rho}_{V_{(\lambda,v)}}(A) = c_{(\lambda,v)}^{-1} \cdot \det(A_{\lambda,v}A_{\lambda,v}^*)\varphi_{V_{(\lambda,v)}^{\mathsf{lin}}}(A - K_{(\lambda,v)}),$$

and we have

$$\frac{\rho_{\hat{\mathcal{V}}}}{\mathrm{NJ}\pi_2}(A,\lambda,v) = \rho_{\mathbb{C}\times\mathbb{S}(\mathbb{C}^n)}(\lambda,v) \cdot \widetilde{\rho}_{V_{(\lambda,v)}}(A).$$

PROOF. It is easy to check that Proposition 6.3 remains valid when replacing $\mathbb{P}(\mathbb{C}^n)$ by $\mathbb{S}(\mathbb{C}^n)$. So for $(A,\lambda,v)\in\hat{\mathcal{V}}$ we get

$$\frac{\mathrm{NJ}\pi_1}{\mathrm{NJ}\pi_2}(A,\lambda,v) = |\det(A_{\lambda,v})|^2 \tag{38}$$

We write $A = M + K_{(\lambda,v)}$ with $M \in V^{\mathsf{lin}}_{(\lambda,v)}$. Combining (38) with (35), we get

$$\begin{aligned}\frac{\rho_{\hat{\mathcal{V}}}}{\mathrm{NJ}\pi_2}(A,\lambda,v) &= \frac{1}{2\pi n}\,\varphi^{\overline{Q},\sigma}_{n\times n}(A)\cdot|\det(A_{\lambda,v})|^2 &(39)\\ &= \frac{1}{2\pi n}\,\varphi_{V^{\mathsf{lin}}_{(\lambda,v)}}(M)\cdot\varphi_{C_{(\lambda,v)}}(K_{(\lambda,v)})\cdot|\det(A_{\lambda,v})|^2. &(40)\end{aligned}$$

For fixed (λ,v) we integrate both sides of this equality over $M\in V^{\mathsf{lin}}_{(\lambda,v)}$. Equality (27) tells us that on the left-hand side we obtain $\rho_{\mathbb{C}\times\mathbb{S}(\mathbb{C}^n)}(\lambda,v)$. On the right-hand side we obtain $\frac{1}{2\pi n}\,c_{(\lambda,v)}\cdot\varphi_{C_{(\lambda,v)}}(K_{(\lambda,v)})$ by the definition (36) of $c_{(\lambda,v)}$. This proves the first equality (37) in the statement.

For the second, we use the definition (28) for the conditional density and (35), (37), (34), and (38) to get

$$\begin{aligned}\widetilde{\rho}_{V_{(\lambda,v)}}(A) &= \frac{\rho_{\hat{\mathcal{V}}}(A,\lambda,v)}{\rho_{\mathbb{C}\times\mathbb{S}(\mathbb{C}^n)}(\lambda,v)\,\mathrm{NJ}\pi_1(A,\lambda,v)}\\ &= \frac{\frac{1}{2\pi n}\,\varphi^{\overline{A},\sigma}_{n\times n}(A)\,\mathrm{NJ}\pi_1(A,\lambda,v)}{\frac{1}{2\pi n}\,c_{(\lambda,v)}\,\varphi_{C_{(\lambda,v)}}(K_{(\lambda,v)})\,\mathrm{NJ}\pi_2(A,\lambda,v)}\\ &= c^{-1}_{(\lambda,v)}\cdot\varphi_{V^{\mathsf{lin}}_{(\lambda,v)}}(M)\cdot|\det(A_{\lambda,v})|^2.\end{aligned}$$

This proves the second inequality. The third equality is a trivial consequence of (39) and the first two assertions of the lemma. □

We can now give the proof of the main result in this section.

PROOF OF THEOREM 2.16. Because of Lemma 6.4(b) (and Remark 6.5) we have

$$\mathop{\mathbb{E}}_{Q\sim\mathcal{N}(\overline{Q},\sigma^2\mathsf{Id})}\Big(\frac{\mu^2_{\mathsf{av}}(Q)}{\|Q\|_F^2}\Big) = \mathop{\mathbb{E}}_{(Q,\lambda,v)\sim\rho_{\hat{\mathcal{V}}}}\Big(\frac{\mu^2(Q,\lambda,v)}{\|Q\|_F^2}\Big).$$

By the definition of the condition number 7 we have

$$\frac{\mu(Q,\lambda,v)}{\|Q\|_F} = \|Q^{-1}_{\lambda,v}\|.$$

Hence

$$\begin{aligned}\mathop{\mathbb{E}}_{(Q,\lambda,v)\sim\rho_{\hat{\mathcal{V}}}}\Big(\frac{\mu^2(Q,\lambda,v)}{\|Q\|_F^2}\Big) &= \mathop{\mathbb{E}}_{(Q,\lambda,v)\sim\rho_{\hat{\mathcal{V}}}}\big(\|Q^{-1}_{\lambda,v}\|^2\big)\\ &= \mathop{\mathbb{E}}_{(\lambda,v)\sim\rho_{\mathbb{C}\times\mathbb{S}(\mathbb{C}^n)}}\Big(\mathop{\mathbb{E}}_{Q\sim\widetilde{\rho}_{V_{(\lambda,v)}}}\big(\|Q^{-1}_{\lambda,v}\|^2\big)\Big) &(41)\end{aligned}$$

the last by equation (29).

Because of unitary invariance, we may assume $v = e_1 = (1, 0, \dots, 0)$. In this case, we have

$$V_{(\lambda,e_1)} = \Big\{ \begin{bmatrix} \lambda & a \\ 0 & B \end{bmatrix} \;\Big|\; a \in \mathbb{C}^{n-1}, B \in \mathbb{C}^{(n-1)\times(n-1)} \Big\}, \quad K_{(\lambda,v)} = \begin{bmatrix} \lambda & 0 \\ 0 & 0 \end{bmatrix},$$

and $V^{\mathsf{lin}}_{(\lambda,e_1)}$ equals the space of matrices of the form $M = \begin{bmatrix} 0 & a \\ 0 & B \end{bmatrix}$. The Gaussian distribution $\varphi_{V^{\mathsf{lin}}_{(\lambda,v)}}$ of M induces an isotropic Gaussian distribution $\rho_\sigma^{(n-1)}$ of $a \in \mathbb{C}^{(n-1)}$ with center $\overline{a}$ and an isotropic Gaussian distribution $\rho_\sigma^{(n-1)^2}$ of $B \in \mathbb{C}^{(n-1)\times(n-1)}$ with center $\overline{B}$.

Decomposing $Q \in V_{(\lambda,e_1)}$ as $Q = M + K_{(\lambda,v)}$, Lemma 6.6 tells us that the conditional density $\widetilde{\rho}_{V_{(\lambda,e_1)}}$ on the fiber $V_{(\lambda,e_1)}$ has the form

$$\begin{aligned} \widetilde{\rho}_{V_{(\lambda,e_1)}}(Q) &= c^{-1}_{(\lambda,e_1)} \cdot |\det(Q_{\lambda,e_1})|^2 \rho_{V^{\mathsf{lin}}_{(\lambda,e_1)}}(M) \\ &= c^{-1}_{(\lambda,e_1)} \cdot |\det(B - \lambda \mathsf{Id})|^2 \, \rho_\sigma^{(n-1)}(a) \, \rho_\sigma^{(n-1)^2}(B). \end{aligned} \tag{42}$$

For the second equality we used that $Q_{\lambda,v}$ is represented by the matrix $B - \lambda\mathsf{Id}$, see (4). By the definition (36) of $c_{(\lambda,e_1)}$,

$$\begin{aligned} c_{(\lambda,e_1)} &= \mathop{\mathbb{E}}_{M \sim \rho_{V^{\mathsf{lin}}_{(\lambda,e_1)}}} |\det(M_{\lambda,v})|^2 \\ &= \int_{\substack{a \in \mathbb{C}^{(n-1)} \\ B \in \mathbb{C}^{(n-1)\times(n-1)}}} |\det(B - \lambda\mathsf{Id})|^2 \, \rho_\sigma^{(n-1)}(a) \, \rho_\sigma^{(n-1)^2}(B) \, da \, dB \\ &= \mathop{\mathbb{E}}_{\widetilde{B} \sim \mathcal{N}(\overline{B} - \lambda\mathsf{Id}, \sigma^2\mathsf{Id})} |\det \widetilde{B}|^2. \end{aligned} \tag{43}$$

It follows from (42) that,

$$\begin{aligned} \mathop{\mathbb{E}}_{Q \sim \widetilde{\rho}_{V_{(\lambda,e_1)}}} \big(\|Q^{-1}_{\lambda,e_1}\|^2\big) &= \int_{\substack{a \in \mathbb{C}^{(n-1)} \\ B \in \mathbb{C}^{(n-1)\times(n-1)}}} \|(B - \lambda\mathsf{Id})^{-1}\|^2 \, c^{-1}_{(\lambda,e_1)} \, |\det(B - \lambda\mathsf{Id})|^2 \\ &\qquad \rho_\sigma^{(n-1)}(a) \, \rho_\sigma^{(n-1)^2}(B) \, da \, dB \\ &= \mathop{\mathbb{E}}_{B \sim \mathcal{N}(\overline{B}, \sigma^2\mathsf{Id})} \|(B - \lambda\mathsf{Id})^{-1}\|^2 \, c^{-1}_{(\lambda,e_1)} \, |\det(B - \lambda\mathsf{Id})|^2 \\ &= \mathop{\mathbb{E}}_{\widetilde{B} \sim \mathcal{N}(\overline{B} - \lambda\mathsf{Id}, \sigma^2\mathsf{Id})} \|\widetilde{B}^{-1}\|^2 \, c^{-1}_{(\lambda,e_1)} \, |\det \widetilde{B}|^2. \end{aligned}$$

The form of this expectation (along with that of $c_{(\lambda,e_1)}$ given in (43)) is exactly the one in the hypothesis of [7, Proposition 4.22], a result then ensuring that

$$\mathop{\mathbb{E}}_{Q \sim \widetilde{\rho}_{V_{(\lambda,e_1)}}} \big(\|Q^{-1}_{\lambda,e_1}\|^2\big) \leq \frac{en}{2\sigma^2}.$$

The fact that this last bound does not depend on the pair (λ, e_1) allows to quickly derive our result from (41). □

7 Proof of Theorems 2.22 and 2.23

7.1 A useful change of variables

Since a linear combination (with fixed coefficients) of two Gaussian matrices is Gaussian as well, it is convenient to parameterize the interval $[M, A]$ by a parameter $t \in [0, 1]$ representing a ratio of Euclidean distances (instead of a ratio of angles as τ does). Thus we write, abusing notation, $Q_t = tA + (1-t)M$. For fixed t, as noted before, Q_t follows a Gaussian law. For this new parametrization we have the following result (see [6, Proposition 5.2] for a proof).

Proposition 7.1 *Let $A, M \in \mathbb{C}^{n\times n}$ be $\mathbb{R}$-linearly independent and $\alpha := d_{\mathbb{S}}(A, M)$. The function*

$$\begin{array}{rcl} t : [0,1] & \to & [0,1] \\ \tau & \mapsto & t(\tau) := \dfrac{\|M\|_F}{\|M\|_F + \|A\|_F(\sin\alpha\cot(\tau\alpha) - \cos\alpha)} \end{array}$$

is a bijection satisfying, for every $\tau \in [0, 1]$, that

$$Q_\tau = t(\tau)A + (1 - t(\tau))M.$$

Furthermore, for all $0 \le a \le b \le 1$,

$$d_{\mathbb{S}}(A, M) \int_a^b \mu^2_{\mathsf{av}}(Q_\tau)d\tau \le \|A\|_F \|M\|_F \int_{t(a)}^{t(b)} \frac{\mu^2_{\mathsf{av}}(Q_t)}{\|Q_t\|_F^2}\, dt. \qquad \square$$

7.2 Proof of Theorem 2.22

We want to bound $\mathsf{Avg_Cost}(n) = \mathcal{O}(n^3)\mathsf{Avg_Num_Iter}(n)$. To do so it is enough to bound $\mathsf{Avg_Num_Iter}(n)$. Recall, we have

$$\begin{array}{rcl} \mathsf{Avg_Num_Iter}(n) & = & \displaystyle\mathop{\mathbb{E}}_{A\sim\mathcal{N}(0,\mathsf{Id})} \frac{1}{n}\sum_{j=1}^{n} K(A, M, m_{jj}, e_j) \\ & \le & \displaystyle 1077 \mathop{\mathbb{E}}_{A\sim\mathcal{N}(0,\mathsf{Id})} d_{\mathbb{S}}(M, A)\frac{1}{n}\sum_{j=1}^{n}\int_0^1 \mu^2(Q_\tau, \lambda_\tau^{(j)}, v_\tau^{(j)})d\tau \\ & = & \displaystyle 1077 \int_0^1 \mathop{\mathbb{E}}_{A\sim\mathcal{N}(0,\mathsf{Id})} d_{\mathbb{S}}(M, A)\mu^2_{\mathsf{av}}(Q_\tau)d\tau \qquad (44) \\ & = & \displaystyle 1077 \int_0^1 \mathop{\mathbb{E}}_{A\sim\mathcal{N}_T(0,\mathsf{Id})} d_{\mathbb{S}}(M, A)\mu^2_{\mathsf{av}}(Q_\tau)d\tau \end{array}$$

the second line by Proposition 2.19, the third by the definition of μ_{av}, and the fourth because, for each $\tau \in [0,1]$, the function $A \mapsto d_{\mathbb{S}}(M,A)\mu^2_{\mathsf{av}}(Q_\tau)$ is scale invariant (and we therefore use (9)).

We are left with the task of bounding the last expression in (44). The general idea is to use the change of variables described in §7.1, then use Theorem 2.16 to bound the inner expectation, and finally integrate over $t \in [0,1]$. This direct approach is however infeasible, because the resulting integral in t would be improper. To circumvent this difficulty, the idea (going back to [6]) is simple. For small values of τ the matrix Q_τ is close to M, and therefore, the value of $\mu^2_{\mathsf{av}}(Q_\tau)$ can be bounded by a small multiple of $\mu^2_{\max}(M)$. For the remaining values of τ, the corresponding $t = t(\tau)$ is bounded away from 0, and therefore, so is the variance in the distribution $\mathcal{N}(\overline{Q}_t, t^2\mathsf{Id})$ for Q_t (here $\overline{Q}_t := (1-t)M$). This allows one to control the denominator on the right-hand side of Theorem 2.16 when using this result. Here are the details.

We set $\alpha := d_{\mathbb{S}}(A,M)$, $\varepsilon := 0.12$, $C_\varepsilon := \frac{\varepsilon}{12.5} = 0.0096$, and $\xi := \frac{2C_\varepsilon(1-\varepsilon)}{3\sqrt{6}(1+\varepsilon)^4} \approx 0.001461$, as in the proof of Theorem 2.19, and define

$$T := \sqrt{2}\,n, \quad \delta_0 := \frac{\xi}{\mu^2_{\max}(M)} \qquad \text{and} \qquad t_T := \frac{1}{1+T+1.0000015\,\frac{T}{\delta_0}}.$$

Let $(\lambda^{(1)}, v^{(1)}), \ldots, (\lambda^{(n)}, v^{(n)})$ be the eigenpairs of M and denote by $(Q_\tau, \lambda^{(j)}_\tau, v^{(j)}_\tau)_{\tau\in[0,1]}$ the lifting of $[M,A]$ in $\mathcal{V}$ corresponding to the initial triple $(M, \lambda^{(j)}, v^{(j)})$.

Corollary 5.2 for $i = 0$ implies the following: for all j and all $\tau \le \frac{\xi}{\alpha\,\mu^2(M,\lambda^{(j)},v^{(j)})}$ we have

$$\mu(Q_\tau, \lambda^{(j)}_\tau, v^{(j)}_\tau) \le (1+\varepsilon)\mu(M, \lambda^{(j)}, v^{(j)}) \le (1+\varepsilon)\mu_{\max}(M).$$

In particular, this inequality holds for all j and all $\tau \le \frac{\delta_0}{\alpha}$, and hence for all such τ, we have

$$\mu_{\mathsf{av}}(Q_\tau) \le (1+\varepsilon)\mu_{\max}(M). \tag{45}$$

Splitting the integral in the last last expression in (44) at $\tau_0 := \min\{1, \frac{\delta_0}{d_{\mathbb{S}}(A,M)}\}$, we obtain

$$\begin{aligned}\int_0^1 \mathop{\mathbb{E}}_{A\sim\mathcal{N}_T(0,\mathsf{Id})} \alpha\mu^2_{\mathsf{av}}(Q_\tau)d\tau \;&=\; \mathop{\mathbb{E}}_{A\sim N_T(0,\mathsf{Id})}\Big(\alpha\int_0^{\tau_0}\mu^2_{\mathsf{av}}(Q_\tau)\,d\tau\Big)\\ &\quad+\mathop{\mathbb{E}}_{A\sim N_T(0,\mathsf{Id})}\Big(\alpha\int_{\tau_0}^1\mu^2_{\mathsf{av}}(Q_\tau)\,d\tau\Big).\end{aligned} \tag{46}$$

Using (45) we bound the first term on the right-hand side as follows:

$$\mathop{\mathbb{E}}_{A\sim N_T(0,\mathsf{Id})}\Big(\alpha\int_0^{\tau_0}\mu^2_{\mathsf{av}}(Q_\tau)\,d\tau\Big) \le \delta_0\,(1+\varepsilon)^2\,\mu^2_{\max}(M) = (1+\varepsilon)^2\xi \le 0.002.$$

For bounding the second term, we assume without loss of generality that $\tau_0 \leq 1$. We then use Proposition 7.1 to obtain that for a fixed A (recall $\|M\|_F = 1$)

$$\alpha \int_{\tau_0}^{1} \mu_{\mathsf{av}}^2(Q_\tau)\, d\tau \leq \int_{t_0}^{1} \|A\|_F \frac{\mu_{\mathsf{av}}^2(Q_t)}{\|Q_t\|_F^2}\, dt \tag{47}$$

with

$$t_0 = \frac{1}{1 + \|A\|(\sin\alpha \cot\delta_0 - \cos\alpha)}.$$

Now note that $\|A\|_F \leq T$, since we draw A from $N_T(0, \mathsf{Id})$. This allows us to bound t_0 from below by a quantity independent of A. Indeed, we first note that

$$0 \leq \sin\alpha\cot\delta_0 - \cos\alpha \leq \frac{1}{\sin\delta_0} - \cos\alpha \leq \frac{1}{\sin\delta_0} + 1,$$

and moreover, $\sin\delta_0 \geq 0.9999985 \cdot \delta_0$, since $\delta_0 \leq 2\xi \leq 0.002922$ (Lemma 2.6). We can now use that $\|A\|_F \leq T$ and bound t_0 as

$$t_0 \geq \frac{1}{1 + T + \frac{T}{\sin\delta_0}} \geq \frac{1}{1 + T + 1.0000015\,\frac{T}{\delta_0}} = t_T.$$

We next use this bound, together with (47), and bound the second term in (46):

$$\begin{aligned}
\mathop{\mathbb{E}}_{A\sim N_T(0,\mathsf{Id})} \Big(\alpha \int_{\tau_0}^{1} \mu_{\mathsf{av}}^2(Q_\tau)\, d\tau\Big) \;&\leq\; \mathop{\mathbb{E}}_{A\sim N_T(0,\mathsf{Id})} \Big(T \int_{t_T}^{1} \frac{\mu_{\mathsf{av}}^2(Q_t)}{\|Q_t\|_F^2}\, dt\Big) \\
&= T \int_{t_T}^{1} \mathop{\mathbb{E}}_{A\sim N_T(0,\mathsf{Id})} \left(\frac{\mu_{\mathsf{av}}^2(Q_t)}{\|Q_t\|_F^2}\right) dt \;\leq\; \frac{T}{P_{T,1}} \int_{t_T}^{1} \mathop{\mathbb{E}}_{A\sim \mathcal{N}(0,\mathsf{Id})} \left(\frac{\mu_{\mathsf{av}}^2(Q_t)}{\|Q_t\|_F^2}\right) dt.
\end{aligned}$$

Observing that for fixed t and when A is distributed following $\mathcal{N}(0, \mathsf{Id})$, the variable $Q_t = (1-t)M + tA$ follows the Gaussian $\mathcal{N}(\overline{Q}_t, t^2\mathsf{Id})$, we deduce (recall $T = \sqrt{2}n$ and $P_{T,1} \geq 1/2$ by Lemma 2.18)

$$\int_0^1 \mathop{\mathbb{E}}_{A\sim\mathcal{N}_T(0,\mathsf{Id})} \alpha\mu_{\mathsf{av}}^2(Q_\tau) d\tau \leq 0.002 + 2\sqrt{2}\, n \int_{t_T}^{1} \mathop{\mathbb{E}}_{Q_t\sim\mathcal{N}(\overline{Q}_t, t^2\mathsf{Id})} \left(\frac{\mu_{\mathsf{av}}^2(Q_t)}{\|Q_t\|_F^2}\right) dt.$$

To bound the integral in the right-hand side we apply Theorem 2.16 and obtain

$$\begin{aligned}
\int_{t_T}^{1} \mathop{\mathbb{E}}_{Q_t\sim\mathcal{N}(\overline{Q}_t, t^2\mathsf{Id})} \Big(\frac{\mu_{\mathsf{av}}^2(Q_t)}{\|Q_t\|_F^2}\Big)\, dt \;&\leq\; \int_{t_T}^{1} \frac{en}{2t^2}\, dt = \frac{en}{2}\Big(\frac{1}{t_T} - 1\Big) \;=\; \frac{enT}{2}\Big(1 + \frac{1.0000015}{\delta_0}\Big) \\
&=\; \frac{en^2}{\sqrt{2}}\Big(1 + \frac{1.0000015\,\mu_{\max}^2(M)}{\xi}\Big) \;=\; \mathcal{O}(n^4)
\end{aligned}$$

the last by Lemma 2.20. We conclude that

$$\mathop{\mathbb{E}}_{A\sim N_T(0,\mathsf{Id})} \Big(\alpha \int_{\tau_0}^{1} \mu_{\mathsf{av}}^2(Q_\tau)\, d\tau\Big) \;\leq\; 0.002 + \mathcal{O}(n^5)$$

and hence, that $\mathsf{Avg_Num_Iter}(n) = \mathcal{O}(n^5)$, and that $\mathsf{Avg_Cost}(n) = \mathcal{O}(n^8)$.

We finally prove the smoothed analysis bounds. Reasoning as in (44) we see that the smoothed number of iterations $\mathsf{Smoothed_Num_Iter}(n, \sigma)$ of $\mathsf{Single_Eigenpair}$ satisfies

$$\mathsf{Smoothed_Num_Iter}(n,\sigma) \;\leq\; 1077 \sup_{\overline{A}\in\mathbb{S}(\mathbb{C}^{n\times n})} \int_0^1 \mathop{\mathbb{E}}_{A\sim\mathcal{N}_T(\overline{A},\sigma^2\mathsf{Id})} \alpha\,\mu_{\mathsf{av}}^2(Q_\tau)d\tau. \quad (48)$$

We deal with the integral as above to obtain

$$\int_0^1 \mathop{\mathbb{E}}_{A\sim\mathcal{N}_T(\overline{A},\sigma^2\mathsf{Id})} \alpha\mu_{\mathsf{av}}^2(Q_\tau)d\tau \leq 0.002 + 2(\sqrt{2}n+1)\int_{t_T}^1 \mathop{\mathbb{E}}_{Q_t\sim\mathcal{N}(\overline{Q}_t,\sigma^2t^2\mathsf{Id})} \left(\frac{\mu_{\mathsf{av}}^2(Q_t)}{\|Q_t\|_F^2}\right) dt$$

where we used $\|A\|_F \leq T + \|\overline{A}\|_F = \sqrt{2}n+1$ and we now have $\overline{Q} := (1-t)M + t\overline{A}$. The rest of the reasoning follows by noting that the term $\frac{1}{\sigma^2}$ can be factored out the integral and that the resulting bound for the integral in (48), $\mathcal{O}(\frac{n^5}{\sigma^2})$, is independent of $\overline{A}$. □

7.3 Proof of Theorem 2.23

It is immediate from the fact that, for any $A \in \mathbb{C}^{n\times n}$, the number of iterations performed by $\mathsf{All_Eigenpairs}$—to compute the n eigenpairs of A— is n times the (expected) number of iterations performed by $\mathsf{Single_Eigenpair}$—to compute one such eigenpair. □

References

[1] D. Armentano. Complexity of Path-Following Methods for the Eigenvalue Problem. *Found. Comput. Math.*, 14(2):185–236, 2014.

[2] D. Armentano and F. Cucker. A randomized homotopy for the hermitian eigenpair problem. To appear at *Found. Comput. Math.*, 2014.

[3] C. Beltrán and L.M. Pardo. Smale's 17th problem: average polynomial time to compute affine and projective solutions. *J. Amer. Math. Soc.*, 22(2):363–385, 2009.

[4] C. Beltrán and L.M. Pardo. Fast linear homotopy to find approximate zeros of polynomial systems. *Found. Comput. Math.*, 11(1):95–129, 2011.

[5] I. Briquel, F. Cucker, J. Peña, and V. Roshchina. Fast computation of zeros of polynomial systems with bounded degree under finite-precision. *Math. of Computation*, 83:1279–1317, 2014.

[6] P. Bürgisser and F. Cucker. On a problem posed by Steve Smale. *Annals of Mathematics*, 174:1785–1836, 2011.

[7] P. Bürgisser and F. Cucker. *Condition*, volume 349 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 2013.

[8] K.P. Choi. On the medians of gamma distributions and an equation of Ramanujan. *Proc. Amer. Math. Soc.*, 121(1):245–251, 1994.

[9] J.W. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.

[10] E. Lahaye. Une méthode de resolution d'une categorie d'equations transcendantes. *C. R. Acad. Sci. Paris*, 198:1840–1842, 1934.

[11] T.Y. Li. Numerical solution of polynomial systems by homotopy continuation methods. In Ph. Ciarlet and F. Cucker, editors, *Handbook of Numerical Analysis*, volume XI, pages 209–304. North-Holland, Amsterdam, 2003.

[12] M. Shub and S. Smale. Complexity of Bézout's Theorem I: geometric aspects. *Journal of the Amer. Math. Soc.*, 6:459–501, 1993.

[13] M. Shub and S. Smale. Complexity of Bézout's Theorem II: volumes and probabilities. In F. Eyssette and A. Galligo, editors, *Computational Algebraic Geometry*, volume 109 of *Progress in Mathematics*, pages 267–285. Birkhäuser, 1993.

[14] M. Shub and S. Smale. Complexity of Bézout's Theorem III: condition number and packing. *Journal of Complexity*, 9:4–14, 1993.

[15] M. Shub and S. Smale. Complexity of Bézout's Theorem IV: probability of success; extensions. *SIAM J. of Numer. Anal.*, 33:128–148, 1996.

[16] M. Shub and S. Smale. Complexity of Bézout's Theorem V: polynomial time. *Theoretical Computer Science*, 133:141–164, 1994.

[17] S. Smale. Newton's method estimates from data at one point. In R. Ewing, K. Gross, and C. Martin, editors, *The Merging of Disciplines: New Directions in Pure, Applied, and Computational Mathematics*. Springer-Verlag, 1986.

[18] S. Smale. Complexity theory and numerical analysis. In A. Iserles, editor, *Acta Numerica*, pages 523–551. Cambridge University Press, 1997.

[19] S. Smale. Mathematical problems for the next century. *Mathematical Intelligencer*, 20:7–15, 1998.

[20] D.A. Spielman and S.-H. Teng. Smoothed analysis of algorithms. In *Proceedings of the International Congress of Mathematicians*, volume I, pages 597–606, 2002.

[21] D.A. Spielman and S.-H. Teng. Smoothed analysis: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51(3):385–463, 2004.

[22] D.A. Spielman and S.-H. Teng. Smoothed analysis: An attempt to explain the behavior of algorithms in practice. *Communications of the ACM*, 52(10):77–84, 2009.

[23] J. von Neumann and H.H. Goldstine. Numerical inverting matrices of high order, II. *Proceedings of the Amer. Math. Soc.*, 2:188–202, 1951.